

PLAUSIBLE HYPERCOMPUTABILITY

ARAN NAYEBI

ABSTRACT. For over a decade, the hypercomputation movement has produced computational models that in theory solve the algorithmically unsolvable, but they are not physically realizable according to currently accepted physical theories. While opponents to the hypercomputation movement provide arguments against the physical realizability of specific models in order to demonstrate this, these arguments lack the generality to be a satisfactory justification against the construction of *any* information-processing machine that computes beyond the universal Turing machine. To this end, I present a more mathematically concrete challenge to hypercomputability, and will show that one is immediately led into physical impossibilities, thereby demonstrating the infeasibility of hypercomputers more generally. This gives impetus to propose and justify a new direction that hypercomputability should take in order to be physically possible, at least in principle. Instead of attempting to rely on infinities such as idealized limits of infinite time or numerical precision, or some other physically unattainable source, the field should be geared toward extending the classical paradigm to better encapsulate modern computational problems that are not well-expressed/modeled by the closed-system paradigm of the Turing machine. I present the first steps toward this goal by considering contemporary computational problems dealing with intractability and issues surrounding cyber-physical systems, and argue that a “plausible hypercomputability” should focus on these issues in order to be practically viable.

CONTENTS

1. Introduction	2
2. Models of hypercomputation	3
2.1. Artificial Recurrent Neural Networks (ARNNs)	3
2.2. Elementary models of hypercomputation	4
2.3. Quantum mechanical and relativistic hypercomputation	6
2.4. Biophysical hypercomputation	7
3. A general negative answer to physical hypercomputation	7
3.1. A deeper analysis of Turing’s and Gandy’s arguments	8
3.2. Addressing arguments against Gandy’s thesis	10
3.3. The “test” for hypercomputation: black box approach	16
4. The practical limitations of Turing-equivalent machines: A proposal	18
4.1. Computational theory vs. Computational practice	18
4.2. Towards a plausible hypercomputability	20
5. Conclusions	25
Acknowledgements	26
References	27

1. INTRODUCTION

The seminal work of Gödel, Church, Turing, and Post in the 1930s regarding the analysis of algorithmic process provided a framework for future theoretical and practical investigations in computer science. Letting “effective calculability” refer to an informal notion, and “computable” to mean “computable by a Turing machine”¹, Church’s [9] and Turing’s [100] analyses of effective calculability ultimately led to

Church’s Thesis. *A function of positive integers is effectively calculable iff it is recursive.*

Turing’s Thesis. *What is effectively calculable is computable.*

When restricted to functions of positive integers, both Church’s thesis and Turing’s thesis are equivalent notions, hence the term “Church-Turing thesis”, which was introduced by Kleene [53] in 1967. Kleene [52] provides extensive evidence in support of the Church-Turing thesis; one important justification being that every known effectively calculable function has been shown to be computable. Turing had in mind calculation by an idealized human being who proceeds mechanically with pencil and paper. A machine in the sense of Turing is essentially a finite state device operating on a one-way finite (but potentially infinite) linear tape divided into squares, meant to be deterministic and sequential, comparable to “a man in the process of computing a real number [in that it is] only capable of a finite number of conditions” [100, pg. 231]. Although the Turing machine was meant as a gedankenexperiment for proving the algorithmic unsolvability of the *Entscheidungsproblem*, the early computers of the 1950s were, in some sense, physical realizations of this idealization. Thus, the algorithmically unsolvable problems of the theoretical realm were generally believed to remain unsolvable in the practical realm as well, namely, one could not build a physical device that could solve the algorithmically unsolvable. Though theoretical devices that can compute more than the universal Turing machine have been around since 1939 (starting with Turing’s oracle machine [101]), the “hypercomputation” movement has been relatively recent.

A hypercomputer is any information-processing machine, notional or real, that is able to achieve more than the traditional human clerk working by rote. Hypercomputers compute functions or numbers, or more generally solve problems or carry out tasks, that lie beyond the reach of the universal Turing machine of 1936. Copeland [17, pg. 462]

One aim of hypercomputation has been to argue for the existence of *physical* devices that can compute non-computable functions or problems. The devices presented so far do not work in practice because these hypercomputational models rely on embedded infinities to perform computations, or on some black box oracle that computes a non-recursive function, without much (or any) description of how this oracle explicitly goes about its computations. It is common for opponents of the hypercomputation movement to point to embedded infinities in order to argue against the physical realizability of *specific* models, but the issue with hypercomputation is far deeper and more general than that.

I will argue that the aims of the hypercomputation movement are not well-defined. Thus, what is needed is a deeper analysis of Turing, and I will argue that any sort of “plausible hypercomputability” would have to explain what it is that it expects to compute that is outside the scope of Turing [100] in at least as mathematically precise a manner as Turing did. I will demonstrate that attempts to make the current aims of the hypercomputation movement mathematically precise immediately lead to physical impossibilities given current physical theories, and this gives us impetus to shift the current focus of hypercomputation. Instead of asserting that non-computable functions, tasks, etc. are physically computable, a more reasonable starting point for a more plausible extension to the classical paradigm should be focused on modeling processes that are not adequately encapsulated by the Turing machine as a closed system; as guiding examples, I consider contemporary computational problems dealing with issues surrounding cyber-physical systems as well as intractability – the question of whether or not these processes allow one to physically compute the non-computable is not considered.

¹In places where clarification may be necessary, we will use the equivalent term “Turing-computable”. When referring to the non-Turing computable, we will simply use the standard term “non-computable”.

2. MODELS OF HYPERCOMPUTATION

Can one, given current physical theories, build a machine that solves, say, the halting problem or some other function not computable by the universal Turing machine? Since there have been many such proposed models (and at least speculation of physical processes that cannot be simulated by Turing machine has spanned over four decades²), the latter question can be reformulated as an inquiry into the physical realizability of these models. The most well-known critique of hypercomputation and negative answer to the above question is provided by Martin Davis [25][26][27]. His arguments (and those of other authors) against these models primarily fall under what I will call the “Argument by Infinity” and the “Argument by Physical Vagueness”.

Argument by Infinity. *A model that involves embedded infinities is not physically realizable.*

Argument by Physical Vagueness. *It is an open question/unclear how to take advantage of certain physical phenomena required by a particular model in order to compute the non-computable (not to mention that there may be question about the existence of the physical phenomena in question), or how to construct the particular hypercomputer based on the physical theory (or theories) that it is claimed to be consistent with.*

Note that it is impractical to categorize all arguments against the existence of hypercomputers due to the rather large literature on the subject, but these two cover most (if not all) of them. Of course, the Arguments by Infinity and Physical Vagueness do not just apply to models of hypercomputation but to *any* computational model that involves infinities or is physically vague. The Argument by Physical Vagueness is quite often the consequence of a model’s attempt to utilize an infinite process to compute the non-computable; thus, it is not unusual that the two arguments go together. Furthermore, it should be noted that the Argument by Physical Vagueness *alone* is weaker than the Argument by Infinity, because unlike the latter, the former does not rule out a model (since we are operating under the plain assumption that the physically constructible is finite³), but just makes its existence in the near future highly unlikely.

As one can see, the Argument by Infinity and the Argument by Physical Vagueness are *model-specific*, namely, one applies it to a particular model to rule out its physical plausibility. They by no means provide a satisfactory argument against the physical construction of *any* information processing machine that computes the non-computable given current physical theories, but the fact that hypercomputational models presented so far have not been shown to be physically feasible in part mostly due to the Argument by Infinity and the Argument by Physical Vagueness gives one motivation to pursue more general arguments against hypercomputation, that I will do in §3.

In this section, my focus will be to survey several prominent models of hypercomputation and apply the aforementioned model-specific arguments against them. This survey should not be considered a comprehensive enumeration of the hypercomputational models presented thus far in the literature, but is intended to provide the reader with a flavor of what hypercomputation is and how the arguments against it are employed. A reader already familiar with the hypercomputation movement may wish to begin with §3 and return to §2 as needed.

2.1. Artificial Recurrent Neural Networks (ARNNs). The brain is an incredibly complicated network of neurons, and it is not surprising to see a vast array of computational models based on nervous systems such as the human brain. The notion that cognitive processes can be modeled by a finite interconnection of logical devices has been known since the 1940s and 1950s, beginning with the 1943 paper of McCulloch and Pitts [61]. In 1995, Siegelmann [90], proposed that her Artificial Recurrent Neural Network (ARNN) could compute the non-computable. A detailed monograph [91] was subsequently published purporting the same ideas. From a broad perspective, the essential notion is that hypercomputation can be performed by

²Copeland [13, pp. 130-1] provides an excellent list of papers concerned that present these speculations prior to the coining of the term “hypercomputation” by Copeland and Proudfoot [14] in 1999.

³Though one can never exclude the possibility that a groundbreaking new physical theory could provide for the physical implementation of the infinite; however, so far no such theory is even closely there.

using infinite precision real weights for the synapses, which falls victim to the Argument by Infinity. In 2001, Costa, Rodrigues, and Siegelmann [20] explored the use of formal methods in verifying temporal properties of ARNNs, and when faced with the problem of proving properties about ARNNs with an infinite state space, they remark “We don’t have the answer yet and the work on finite ARNNs has not given us much clues. We believe we should recover the work on finite ARNNs and try to improve the results obtained” [20, pg. 165]. Davis [25] is the most well-known critic of the practicality of ARNNs as a means of hypercomputation, and his assertions (which shall not be reproduced here), lead to the conclusion that the only way in which Siegelmann’s neural nets can recognize a non-computable language is to use non-computable weights. Hence, the neural nets can only compute the non-computable if the weights are already non-computable, again pointing to its physical infeasibility by the Argument by Physical Vagueness.

2.2. Elementary models of hypercomputation. I shall now present several elementary models of hypercomputation, namely those that are relatively well-known and have been around before or since the turn of the century and are usually the basis for most other proposed hypercomputational models. Copeland and Proudfoot [14] were the first to introduce the term “hypercomputation” in 1999 (Siegelmann had used another, equivalent, moniker for her ARNNs in 1995, “super-Turing computation”), using Turing’s doctoral dissertation [101] as an initial basis for this concept. Turing describes the “oracle machine” (or *o*-machine) as follows

Let us suppose that we are supplied with some unspecified means of solving number-theoretic problems; a kind of oracle as it were. We shall not go any further into the nature of this oracle apart from saying that it cannot be a machine. With the help of an oracle we could form a new kind of machine (call them *o*-machines), having as one of its processes that of solving a given number-theoretic problem. Turing [101, pp. 172-3]

Though the description of the *o*-machine constituted a small part of Turing’s doctoral thesis, the introduction of these *o*-machines was the first important step towards classifying degrees of unsolvability from the results of effective unsolvability previously established in the early 1930s for various mathematical problems (cf. Feferman [38]). As noted in Soare [94, §4.2], there are many equivalent modern definitions of an *o*-machine which produce the same class of functionals. Roughly speaking however, an *o*-machine is a Turing machine which can communicate with a black-box oracle *during* computation. This oracle can be thought of as capable of answering a series of questions, represented as some subset of the natural numbers, and can decide the membership of this set. The relevance here to the hypercomputation movement is that an oracle can constitute the “non-classical” part of a particular *o*-machine which returns the values of the “halting function” $H(m, n)$ defined as such

$$H(m, n) := \begin{cases} 1 & \text{if and only if the } m\text{-th Turing machine eventually halts on input } n; \\ 0 & \text{if otherwise} \end{cases}.$$

But the question is, can one physically construct an oracle? Copeland [13] speculates that if physical processes do exist that cannot be simulated by Turing machine then “perhaps future engineers will use them to implement the non-classical part of some *O*-machine” [13, pg. 131]. Even beginning with Turing, who did not intend for the *o*-machine to be built and whose original definition of the oracle was an “unspecified” means of solving number-theoretic problems, it is not surprising that Copeland’s latter speculation is inherently vague (since it is unspecified how such a physical process, even if it exists, could be used to build an oracle), and therefore clearly susceptible to the Argument by Physical Vagueness. A year later, Copeland and Proudfoot [14] provide some clarification into specifying how an oracle could be made. Namely, they propose that information can be represented as the successive digits of an infinite precision real number, but this would be ruled out as physically plausible by the Argument by Infinity, as Davis [27] points out.

Copeland and Sylvan [15] introduced the coupled Turing machine, a Turing machine with one or more

input channels that provide input to the machine as the machine operates.⁴ The device may also possess one or more output channels. The machine's input, for instance, could take the form of a symbol from its alphabet that is written on the first square of the machine's tape, and the specific sequence of input determines the functions that the coupled Turing machine can compute. It is therefore easy to see that what is required for the coupled Turing machine to compute the non-computable would be a non-recursive sequence. In the same paper, Copeland and Sylvan also discuss asynchronous networks of Turing machines, where each Turing machine has a timing functions $\Delta_k : \mathbb{N} \rightarrow \mathbb{N}$ that represent the units of time between its execution of its n -th action and its $(n + 1)$ -th action. A non-recursive source would be needed to be used by the timing function in order for the asynchronous network of Turing machines to compute the non-computable. Another form of hypercomputation, known as the infinite state Turing machine, is a Turing machine that, as the name implies, has an infinite set of states. One property of this type of hypercomputer is given an infinite program but in any single computation only a finite, unbounded amount is utilized. In the early 20th century, Blake [5], Weyl [108], and Russell [80] independently observed that a process that performs its first step in one unit of time and each subsequent step in half the time as the preceding step, can complete an infinity of steps in two units of time. A universal Turing machine working in this manner is called an accelerated Turing machine. This device is capable of deciding the halting problem; not to mention that the definition of the Turing machine does not specify how long an individual step may take during computation, so the argument is that the accelerated Turing machine does not conflict mathematically with the classical notion (cf. Calude and Păun [6]). To avoid too much repetition, a detailed description of the aforementioned elementary hypercomputers and other elementary hypercomputers not mentioned here is provided by Ord [71, §3].

For instance, regarding the case of coupled Turing machines, asynchronous networks of Turing machines, and any other form of hypercomputation that consists of a classical part and a non-recursive information source (ω -machines included), the central requirement for their physical realizability is, of course, a non-recursive information source. One way to do this would be to use a real-valued quantity, such as a particle with a non-recursive charge, but the physical drawback that comes with this method is the requirement of arbitrary precision measurement, which is impossible to attain physically. Another approach of accessing a non-recursive information source may be by way of the observation of a non-recursive physical process. Particle decay is examined as a possible approach by Ord, but it is readily pointed out that while it may be physically harnessable, it is "completely useless for computing definite mathematical results that could not be computed with a Turing machine" [71, pg. 26]. Thus, the question arises as to whether there exist other non-recursive information sources that may be *mathematically* harnessable and useful. One possibility is if the bits of a constant such as π or τ were given by some natural process, but there is little possibility that such a natural process exists due to the Argument by Infinity. Regarding the case of infinite state Turing machines, the physical requirement is an explicit infinite specification. While certain infinite bitstrings (such as the binary expansion of π) can be understood, to gain the true power of infinite state Turing machines, we would need to provide arbitrary infinite specifications. The ability to provide arbitrary infinite specifications is not physically attainable again due to the Argument by Infinity.

Regarding the case of accelerated Turing machines and infinite time Turing machines, an infinite amount of memory is required. Accelerated Turing machines are less dependent on their infinite memory since this requirement is only needed when they do not halt. Of course, in terms of the physical realizability of the aforementioned devices, the infinite memory requirement is rather implausible (clearly since it is victim to the Argument by Infinity). It has been suggested by Calude and Pavlov [7] that the storage of infinite bitstrings can be done via infinite quantum superpositions, but "it is not at all clear that infinite memory is a physical possibility" [71, pg. 25].

Regarding hypercomputers that require infinite computation to compute the non-computable, such as the accelerated Turing machine, which attempts to perform an infinite number of steps in computations in a finite amount of time, one immediately meets with the physically implausible because the head would need to travel faster than the speed of light. Calude and Pavlov [7] assert that infinite computation may

⁴Note that the Turing machine does not interact with its environment during computation.

be achieved via quantum mechanical methods by performing the steps in parallel. Nevertheless, as Shagrir [82] argues, while an infinite time Turing machine computes the halting function $H(m, n)$, accelerated Turing machines that have exactly the same computational structure as Turing machines compute exactly the computable functions. Therefore, while performing supertasks enables accelerated Turing machines to complete infinitely many steps in a finite interval of time, it does not enable them to compute what Turing machines cannot compute. I will discuss the matter of whether accelerating Turing machines can compute the non-computable in more detail in §3.2.4, when addressing the more general possibility of supertasks as hypercomputation.

2.3. Quantum mechanical and relativistic hypercomputation. Kieu [49] [50] proposes a quantum algorithm by way of quantum adiabatic evolution for solving Hilbert’s tenth problem (known to be unsolvable following the work of Davis, Putnam, and Robinson [24] and Matiyasevich [59]). Andrew Hodges [44] criticizes the physical realizability of Kieu’s algorithm (in particular referencing the paper of Smith [92] which presents three counterexamples to Kieu’s argument, particularly concerning nonground states arising with a high probability greater than 99.999%); Kieu addresses these criticisms, ultimately concluding that:

The only central argument against physical implementations of the algorithm is based on an assumption that its Hamiltonians cannot be effectively constructed due to a lack of infinite precision. To show that this may not be a forgone conclusion, we draw attention to the virtually unknown fact that, on the contrary, simple instances of Diophantine equations with apparently *infinitely precise* integer coefficients have *already* been realised in certain experiments known as quantum phase transitions! We also speculate on how [the] central limit theorem might be of some help in the effective implementation of the required Hamiltonians. Kieu [51, pg. 7]

The physical realizability of Kieu’s assertions are extremely unlikely (see, for example, Hagar and Korolev [42]), though he proposes the usage of “occupation numbers” instead of energy. However, as Davis mentions

Kieu’s proposed solution has been thoroughly criticized by people who understand quantum mechanics much better than I do, and I don’t propose to say a great deal about it. But there is one point that I do want to make because it illustrates how blithely the hypercomputationists manage to be blind to the significance of infinity. Kieu suggests that one way his method could be used to determine whether the given equation has a solution is to observe, when the process concludes, a tuple of positive integers, the “occupation numbers”. These are then to be substituted into the given equation. Either they satisfy the equation, in which case we know there is a solution, or they do not, in which case Kieu assures us, the equation has no solutions. Now, evidently there are equations that have positive integer solutions, but for which the least such solution is enormous, for example so large, that to write the numbers in decimal notation would require a space larger than the diameter of our galaxy! In what sense could such numbers be read off a piece of equipment occupying a small part of our small planet? And how can we suppose that it would be feasible to substitute numbers of such magnitude into an equation and to carry out the arithmetic needed to determine whether they satisfy the equation? Davis [26, pg. 130]

Hence, Kieu’s method clearly falls victim to the Argument by Infinity.

It is important to note here that on the other hand, general relativity has also been proposed to compute the non-computable, based on Hogarth’s scheme [45][31][37] regarding the existence of a type of space-time singularity [32], and certain solutions to Einstein’s equations (known as “Malament-Hogarth spacetimes”) support the existence of these singularities. A system based on this singularity solves the halting problem, as is clearly summarized by Stannett [95, §2.5]. Welch [107] provides a far more detailed analysis of the extent

of possible computations in these Malament-Hogarth spacetimes. However, the question of whether such a system is physically possible is a matter of dispute. Earman and Norton [31][32][33] provide a detailed study of this proposal and have identified significant problems with it; hence, Hogarth's scheme is subject to the Argument by Physical Vagueness since there are issues with the physical phenomena in question. Of particular interest is the work of Etesi and N emeti [37] concerning a subclass of Malament-Hogarth spacetimes containing rotating Kerr black holes, thereby permitting an infinite span of time to appear finite to an observer. Even under the assumption that the aforementioned exists in our universe, the question remains as to whether or not it is possible for a device to take advantage of this phenomenon (cf. Davis [27,  31]). Thus, the Argument by Physical Vagueness is again at play here.

2.4. Biophysical hypercomputation. While the previous examples of hypercomputability have their basis in modern physics, a recent paper of Calude and P aun [6] presents the first biological computing agent theoretically having super-Turing capabilities by way of acceleration. Basing their ideas on membrane computing, a branch of molecular computing introduced by P aun [72], Calude and P aun [6] attempt to find classes of membrane systems (called P systems) that theoretically demonstrate deterministic Turing computability and biologically support acceleration. It is important to note that while a biological method of acceleration is an enticing argument for the physical practicality of hypercomputation, regarding the *physical realizability* of the biological methods of acceleration they present, the authors remark in a footnote that “we place ourselves in an idealized framework, where, for instance, there is no limit of acceleration, of decreasing the size, hence we can perform arbitrarily many steps of acceleration” [6, pg. 177, fn. 3]. For now, these results have a more interesting theoretical footing than a physically implementable one, and the hypotheses the authors assume (including, for example, increasing the possibility of reactants, whose collisions are driven by Brownian motion, to react in a shorter time by decreasing the size of “reactors”) have yet to be validated by biophysics, hence the Argument by Physical Vagueness applies. I will discuss P systems in more depth in  4.2.1.

3. A GENERAL NEGATIVE ANSWER TO PHYSICAL HYPERCOMPUTATION

One cannot point to a clear answer to the question, “What is hypercomputation?”. The quote from Copeland presented in  1 is exceptionally vague because when one attempts to clarify what is meant by an “information processing machine” that computes functions or solves problems that “lie beyond the reach of the universal Turing machine”, one is only faced with a plethora of models that supposedly proceed to do this in a variety of different manners, as evidenced by even the brief survey of  2. Not to mention that the common theme between these models is that they are often based on the trivial insight that if a physical process exists that generates a non-computable function (or say, an infinite precision real number), then it could be exploited (somehow) to compute the non-computable. And despite this triviality, there is the further catch that in exploiting such physical processes one falls victim primarily to the Argument by Infinity and/or the Argument by Physical Vagueness.

However, as Stannett comments, “those in favour of hypercomputation can sometimes be too quick to assume that their models do not include hidden infinities, while those against hypercomputation are sometimes too quick to assume that what they consider to be the ingredients of computation are in some sense ‘necessary’ ” [95, pg. 9]. Hence, it may be fair for proponents of hypercomputation to rule my arguments out as careless had I based the physical unrealizability of hypercomputers *solely* on the basis of limited specific cases in  2. Thus, what is developed here is a more mathematically concrete challenge to hypercomputability that is not model-specific, and I will show that one is immediately led into physical impossibilities, therefore demonstrating the effective physical impracticability of hypercomputers more generally.

The best place to start is with a deeper analysis of Turing [100]. Since hypercomputation aims to compute that which lies outside the scope of the universal Turing machine, then what is needed is a mathematically clear conception of that scope.

3.1. A deeper analysis of Turing’s and Gandy’s arguments. Turing’s analysis in [100] is essentially a theory of *human* computation. In §1 of his paper, Turing briefly describes what he means by “computable real number” as one whose decimal expression is “calculable by finite means” [100, pg. 230]. This definition as well as the comparison in §1 to a man in the process of computing a real number to the computation performed by an *a*-machine (which is what has become known as the “Turing machine”⁵) is not justified until §9 of his paper. In this justification, we find he provides three kinds of arguments: “(a) A direct appeal to intuition. (b) A proof of the equivalence of two definitions (in case the new definition has a greater intuitive appeal). (c) Giving examples of large classes of numbers which are computable” [100, pg. 249]. It is evident then that Turing’s emphasis when describing the formal mathematical notion of computability is on how much it “agrees” with one’s intuition of effective calculability, where “effective” is taken to mean that the calculation is deterministic and terminates after a finite amount of time. In fact, this direct appeal to intuition (the first argument) is what is most remembered, in spite of the fact that Turing gave two other arguments in support of his justification, as mentioned above. This intuitive definition is widely accepted, one important reason being its convincing psychological appeal. For instance, Church [10] found Turing’s method regarding the unsolvability of the *Entscheidungsproblem* to be more favorable than his own. Moreover, Gödel praised Turing’s explication stating that it is “absolutely impossible that anybody who understands the question and knows Turing’s definition should decide for a different concept” [105, pg. 84]. Despite this praise, an appeal to intuition is not as mathematically rigorous as possible – what is needed is an axiomatization of Turing’s explication to serve as guiding principles for the inquiry of this section (of course, this axiomatization would be implicit in Turing’s original argument, even though he did not explicitly state it). Certainly Turing’s focus on *human* mechanical calculations was appropriate given the historical context of his work, for any solution of the *Entscheidungsproblem* was to be carried out by us! The question that immediately arises, however, is that even though Turing started by analyzing human computation, why should we follow him in that? The reason is that Turing’s analysis offers a detailed and precise investigation that connected an informally understood concept (that of “effective calculability”) to a mathematically precise notion (that of “computability”). And this analysis, through its use of restrictions, allows us to readily segue into the notion of *machine computation*, which is of prime relevance to the investigative aims of this section regarding a more mathematically precise understanding of hypercomputability.

One primary reason why Turing’s analysis of effective calculability made the previously informally understood concept mathematically rigorous was that he placed constraints on the steps permitted in calculations – these were entirely based on the limitations of the human computing agent. Sieg [89] formalizes these limitations as such:

- (1) **Boundedness:** A computer can immediately recognize only a bounded number of configurations.
- (2) **Locality:** A computer can change only immediately recognizable configurations,

where, under Sieg’s terminology, a “computer” refers to a human computing agent who proceeds mechanically. It is easy to see that processes satisfying the above conditions are computable by Turing machines since the first conditions require that the configurations, on which the operations are performed, have a fixed bounded size; and the second conditions require that the operations take place locally, within a fixed finite neighborhood. The axioms that Sieg [89] derives for Turing computers are motivated by these boundedness and locality conditions; I will not state them here.

It is Gandy [39] who makes explicit the connection between the abstract human being working in a rote manner to the physical construction of a machine. While Turing viewed the aforementioned restrictive provisos due to the memory limitations of computers (stating that the “human memory is necessarily limited” [100, pg. 231]), Gandy, on the other hand, appealed to evident physical limitations of machines as impetus to derive restrictive principles for discrete deterministic mechanical devices. By considering and analyzing cellular automata and parallel computation, he produces four principles that any (non-analog) machine must satisfy and proves that devices satisfying them are Turing-equivalent.

The restrictive conditions that Gandy presents are motivated by the physical considerations of quantum mechanics and special relativity, namely, “that there is a lower bound on the linear dimensions of every

⁵It was first called this by Church in his 1937 review [10] of Turing’s paper.

atomic part of the device and that there is an upper bound (the velocity of light) on the speed of propagation of changes” [39, pg. 126]. Gandy aims to provide an argument for the following thesis:

Thesis M. *What can be calculated by a machine is computable.*

He clarifies that he will consider only discrete mechanical devices, and proceeds to formulate the notion of a discrete mechanical device in terms of four precise axiomatic principles. Devices (synonymous with “machine”) satisfying Gandy’s principles (termed “Gandy machines”) are based on a set-theoretic framework for specification, consisting of classes of “states” with a “transition operation” allowing for a step-wise transition from one state to the next. States are represented by subclasses of hereditarily finite sets (HF) built up from a potentially infinite set of atoms, conceived as the basic components of the machine, closed under isomorphic structures (such subclasses are called “structural classes”). Transitions are given by restricted operations from states to states, termed “structural operations”. Ignoring the technical details of Gandy’s presentation (which are clarified in Sieg and Byrnes [85] and Sieg [88]), we can manage to approximate the four principles as follows:

- (1) **Principle I (The Form of Description):** A Gandy machine M can be described by the pair (S, F) , where S is a structural class, and F is a structural operation from S to S . If $x_0 \in S$ is an initial state of M , then $F(x_0)$, $F(F(x_0))$, \dots are its subsequent states.

The remaining three principles place restrictions on S and F involving certain finiteness or boundedness conditions, and Gandy [39] shows that if any of these principles are significantly weakened in almost any way, then every number-theoretic function becomes calculable.

- (2) **Principle II (The Principle of Limitation of Hierarchy):** The set-theoretic rank of the states is bounded. Roughly speaking, each state of S can be assembled from parts, and those parts can be assembled from other parts, etc., but the complexity of this structure is bounded.
- (3) **Principle III (The Principle of Unique Reassembly):** Each state of S can be assembled from basic parts of bounded size, and there is a unique way of putting these parts together (while model construction-kits aim to satisfy this principle, this principle also takes into consideration parts which do overlap).
- (4) **Principle IV (The Principle of Local Causality):** For any state $x \in S$, $F(x)$ is assembled from bounded (and possibly overlapping) parts of x . This is an abstraction of the aforementioned physical considerations of quantum mechanics and special relativity that Gandy takes into account.

From these four principles, Gandy [39] proves the following

Theorem 1. *What can be calculated by a device satisfying Principles I-IV is computable.*

As a result, he clarifies his arguments for a more definite version of “Thesis M” in the following way:

Thesis P. *A discrete deterministic mechanical device satisfies Principles I-IV.*

In sum, *Thesis P* (or *Gandy’s thesis*) is the assertion that any discrete deterministic mechanical device must satisfy the restrictive principles of local causation and unique assembly, and as Sieg and Byrnes [85] point out, this is analogous to Turing’s Thesis which claims that any human computer, which can be represented by a discrete dynamical system and his operation, must obey appropriate locality conditions. Of course, the important difference between the Gandy machine and the Turing machine is that latter can modify only a bounded part of a state, whereas the former allows state-transitions that result in arbitrarily many bounded parts, since Gandy machines take parallel working into account.

Despite the shortcoming that Principle IV does not apply to machines obeying Newtonian mechanics, since “in these there may be rigid rods of arbitrary lengths and messengers travelling with arbitrary large velocities, so that the distance they can travel in a single step is unbounded” [39, pg. 145], Gandy’s principles are a reasonable (in the sense that they obey the modern physical theories of quantum mechanics and special relativity) and mathematically concrete formulation of physical principles that a real machine⁶ must satisfy.

⁶I am only interested in the consideration of real machines here, though Gandy’s Thesis leaves open to interpretation the possible consideration of both notional and real machines.

Not to mention that since the early twentieth century, it has been known that at extremes of velocity (for instance, near the speed of light) and on incredibly small scales, Newtonian mechanics is not a successful physical theory. Moreover, it is worth mentioning that Sieg [89] obtains an axiomatization of the class of Gandy machines based on Gandy’s principles of local causation and unique assembly, and Sieg’s axioms for discrete dynamical systems (for both human and machine computations) result in the reduction of models of these axioms to Turing machines, and cellular automata and many artificial neural nets, for example, can be shown to satisfy these axioms [89]. This axiomatization thus provides good evidence for Gandy’s principles for mechanisms as sufficiently general, and in fact, in [85] and [88], two formulations of abstract models for parallel computations are given based on Gandy’s [39] precise mathematical framework (I will mention their relation to other paradigms of computing in §4.2.1).

3.2. Addressing arguments against Gandy’s thesis. Copeland and Shagrir [18], however, argue that Gandy’s principles for mechanisms are not general enough, and instead they present “(ideal) physical machines that lie outside the class of ‘Gandy machines’ and that compute (in a broad sense of ‘compute’) non-recursive functions....[including] discrete, deterministic hypercomputers from both Newtonian and relativistic frameworks, as well as non-discrete and non-deterministic hypercomputers” [18, pg. 220]. The focus of this paper is, of course, on the physical possibility of hypercomputation, not the notional one. Although the question of whether discrete deterministic mechanical devices contravening Gandy’s principles are allowed by the physics of the real world is a question that is empirical not logical, which the authors are apt to point out, it is precisely the logical framework of Gandy’s that makes the empirical affirmation of hypercomputation unlikely, as I will show below.

The machines that Copeland and Shagrir present that lie outside the class of Gandy machines comprise none other than a partial survey of hypercomputational models in the literature, many of which were discussed in §2. Before proceeding to evaluate the physical realizability of these machines, it is important that I make further precise my notion of “physical realizability”. The Argument by Infinity and the Argument by Physical Vagueness offered a satisfactory starting point, allowing us to provide arguments against a large number of hypercomputational models. Given that the Church-Turing thesis does not directly apply to what can be computed by physical systems in general (but is concerned with what can be determined by effective procedures), and given that Copeland and Shagrir aim to discuss machines that compute non-recursive functions in a broad sense of “compute” (especially due to the substantial number of hypercomputational models in the literature), the question is: Are there more general constraints on physical computation by which to judge a particular model’s physical realizability?

3.2.1. General constraints on physical computation. Piccinini [75, §2] offers a suitable notion of physical computation, which we can use, and can be formulated as the following “usability constraint” on physical computation:

Usability Constraint. *If a physical process is a computation, it can be used by a finite observer to obtain the desired values of a function.*

This can be further clarified if one clarifies what is meant by “finite observer” and “can be used by a finite observer”. A finite observer, in Piccinini’s sense, is a human being or any other intelligent being of similarly bounded capacity (he actually proposes another broader definition, but the current definition is sufficient for the purposes I intend). As to the issue of what it means for a physical process to be useable by a finite observer, Piccinini formulates the following six sub-constraints, which shall be summarized here:

- (1) **Readable Inputs and Outputs:** The inputs and outputs of a computation must be readable, in the sense that the inputs and outputs can be measured to the desired degree of approximation. For example, having an infinite precision real number be the input and/or output would not be permitted. Furthermore, as Piccinini [75, pg. 742, fn. 13] notes, analog computers do not conflict with this principle as they are designed to yield the values of functions of real variables, which is not quite the same as functions of real-valued quantities.

- (2) **Process-Independent Rule:** In a genuine computation, the problem being solved (equivalently, the function being computed) must be definable *independent* of the process of solving it (or equivalently, computing it). This sub-constraint entails that modeling one physical system using another is not enough for computing. Moreover, any machine that is reputed to solve the halting problem (note the halting function $H(m, n)$ returns either 1 or 0, as mentioned in §2.2) by reliably generating (measurable) values of the function for any argument (of reasonable size) satisfies this second sub-constraint.
- (3) **Repeatability:** For a physical process to be a genuine computation, it must be repeatable by any competent finite observer who intends to obtain its results, in the sense that the same sequence of states can occur either within the same physical system at different times or within two relevantly similar physical systems (e.g. two systems satisfying the same equations).
- (4) **Settability:** An ordinary computing system can compute any value of one or more functions, within its limits (e.g. a universal machine can compute the value of any computable function until it runs out of memory or time). Moreover, a system must be *setttable*, namely, a user can choose which value of the function is going to be generated in a given case.
- (5) **Physical Constructibility:** If a system cannot be physically constructed, it may count as performing notional computations, which are irrelevant for physical purposes. This fifth sub-constraint is certainly *not* satisfied by a machine that is either affected by the Argument by Infinity or the Argument by Physical Vagueness.
- (6) **Reliability:** While the requirement that machines should never break down is unrealistic, it is a requirement that machines be reliable in the sense that they completely operate correctly long enough to yield correct results at least some of the time. A machine that never completes its computation successfully is certainly not worth building.

As noted by Piccinini, the usability constraint is weaker than the verifiability constraint that is discussed and rightfully dismissed by Shagrir and Pitowsky [81, pp. 90-1]. Not to mention that the notion of finite observer that was defined above makes the usability constraint relevant to those who are interested in designing, building, selling, using computers now or in the future. And the above list of six sub-constraints, though certainly not exhaustive, is enough to exemplify the properties that are required to produce a useful physical computation. Moreover, the use of this usability constraint in §4 will further demonstrate its importance. Admittedly, these constraints are preliminary, more work will have to be done in order to sharpen them to truly be “general constraints on physical computation”. For instance, as Sieg points out in personal communication, “On what grounds is a system to be judged as not “physically constructible”?”. Sieg believes that the consequences of his axioms are rather direct: “if a machine or physical process can be viewed as a discrete dynamical system, but performs super-Turing computable tasks, then it must violate one of the physical restrictive conditions [of Gandy, as captured by Sieg’s axioms]”. Whether or not this is the case is a separate (but related) matter that would have to be argued extensively elsewhere. However, for the extensive counterexamples that Copeland and Shagrir present, it turns out that the Argument by Infinity and the Argument by Physical Vagueness suffice for sharpening the fifth sub-constraint. In other words, if a model violates any of these, then it can be said that it is not “physically constructible” (given current physical theories). On the other hand, just because a model is physically constructible does not mean, in our terminology, that it is “physically realizable”. If the machines that Copeland and Shagrir [18] present which fall outside the class of Gandy machines do not satisfy this usability constraint (namely by violating at least one of the six aforementioned sub-constraints), then they are ruled out as “physically realizable” (I will refer to them as “physically unrealizable hypercomputers”, or “physically unrealizable” for short, or any other applicable synonym). I should note that my criterion encompasses Piccinini’s criterion for distinguishing between what he calls “spurious” and “genuine” hypercomputers, where the former fails to satisfy at least one of the first four sub-constraints on physical computation, and the latter satisfies at least the first four sub-constraints.

3.2.2. *Partially random machines.* Setting aside the case of analog hypercomputation for the moment, the first model that Copeland and Shagrir present is the “partially random machine”, introduced by Turing [102] in an unpublished (in his lifetime) manuscript entitled “Intelligent Machinery” as follows:

It is possible to modify the above described types of discrete machines by allowing several alternative operations to be applied at some points, the alternatives to be chosen by a random process. Such a machine will be described as ‘partially random’. If we wish to say definitely that a machine is not of this kind we will describe it as ‘determined’. Sometimes a machine may be strictly speaking determined but appear superficially as if it were partially random. This would occur if for instance the digits of the number π were used to determine the choices of a partially random machine, where previously a dice thrower or electronic equivalent had been used. These machines are known as apparently partially random. Turing [102, pg. 416]

The authors point out that the outputs of a partially random machine may form a non-computable sequence, since as Church [11] argues, if a sequence of digits $a_1, a_2, \dots, a_i, \dots$ is random, then there is no function $g(i) = a_i$ that is computable. However, a random output cannot be used to perform any interesting physical “hypercomputation”, such as solving the halting problem, since given the usability constraint, there are several reasons to deny that a genuine random process R should be considered a *physical* computation. As Piccinini [75, pp. 750-1] mentions, there is no process-independent way to define the function $f : \mathbb{N} \rightarrow \{0, 1\}$ without referring to R itself, thereby violating the second sub-constraint; random processes are not repeatable, thereby violating the third sub-constraint; and lastly, random processes are not settable (which violates the fourth sub-constraint) since the “user” of R cannot select the value of f he or she intends to generate. Of course, random processes can be exploited by a computation, but random processes cannot be used to generate the desired values of a function or solve the desired instances of a general problem, such as the halting problem. Physical unrealizability aside, it is not known, as Copeland and Shagrir mention, whether it is possible for there to be partially random (ideal) machines that offer any interesting counterexamples to Thesis M. Even this possibility is unlikely, for in the case of ideal machines where one *can* refer to the possibility of enumeration of infinite sets and the computation of infinite sequences, Theorem 2 of the classic paper by de Leeuw, et al. [28] essentially states that anything that can be done by a machine with a random element (say, for the sake of example, it is a binary random device with probability p of producing a 1 in its output) can be done by a deterministic machine (given that p is a computable real number⁷).

3.2.3. *Gandy machines that produce a non-computable output.* Of course, the partially random machine is not a Gandy machine; though, the authors also consider machines that do satisfy Gandy’s principles but may produce a non-computable output from computable input. It is worth mentioning though that the underlying assumption being made is that this occurs if these machines are embedded in a universe whose physical laws have non-computability built into them, and therefore the authors place themselves in an imagined universe. Hence, almost immediately, this violates the fifth sub-constraint of physical constructibility. Moreover, the example they consider are asynchronous networks of Turing machines, which I mentioned in §2.2 falls victim to the Argument by Infinity, thereby violating the fifth sub-constraint in another way.

3.2.4. *Supertasks as hypercomputation.* The final set of machines that Copeland and Shagrir discuss are supertasks, in particular the accelerated universal Turing machine (AUTM) and the approach of using supertasks in relativistic space-time, both of which were mentioned to some extent in §2 (particularly §§2.2 and 2.3, respectively). The AUTM is an example of a machine obeying Newtonian mechanics [16] and hence falls outside of the class of Gandy machines since it does not obey Principle IV. Specifically, it violates

⁷In fact, given Turing’s description of partially random machines, it seems that the means by which the machine is to make its choices in order to appear it has a random element makes the restriction that p be a computable real number rather reasonable.

Gandy’s requirement that “that there is an upper bound (the velocity of light) on the speed of propagation of changes” [39, pg. 126]. The authors also include Davies’ [22] proposal for the construction of an infinitely accelerating Turing machine which involves an infinite series of machines that, after performing some computational operations, builds a smaller and faster copy of itself, and so on, infinitely so. Again this proposal is consistent with Newtonian mechanics, for Davies’ goal was to produce machines similar to Babbage’s Analytical Engine, “consistent with any physics known in the year 1850” [22, pg. 672]. Moreover, although this does not break the upper-bound on the velocity of light, it still does not obey Principle IV as it depends on an absence of a lower bound on the size of the atoms. Not only is this latter proposal evidently susceptible to the Argument by Infinity, but a relevant question is why would one want to conform to what was known about physics around 1850? Precisely where infinities arise in Newtonian frameworks for hypercomputation and where Newtonian mechanics breaks down are at extremes such as the velocity of light or on the atomic scale. Even as Davies admits, the quantum mechanical nature of matter makes his construction physically impossible. Thus, I shift my focus to hypercomputation by way of modern physical theories, and inquire as to whether or not this is physically realizable. The supertasks in relativistic space-time that Copeland and Shagrir present were already introduced in §2.3. It is worth noting that Piccinini considers these PMH (Pitowsky-Malament-Hogarth) relativistic hypercomputers to be “genuine” hypercomputers, as they satisfy sub-constraints 1-4 [75, §4.2], but at the cost of their physical constructibility (sub-constraint 5) and reliability (sub-constraint 6) being unlikely. For as was noted in §2.3, the relativistic hypercomputers proposed so far are susceptible to the Argument by Physical Vagueness.

But one can say more general things about supertasks as hypercomputation. In §2.2, I mentioned Shagrir’s [82] clear assertions that simply because the AUTM performs a supertask, does not mean it can compute what Turing machines cannot compute; instead, what matters is the *computational structure* involved. Since the AUTM and the Turing machine have the same structure, their computational power is equivalent. However, this does not mean that supertask machines cannot be hypercomputers (but their hypercomputational capabilities have less to do with their performance of a supertask and primarily has to do with their structure differing from that of the Turing machine). As examples of this, in a later paper, Shagrir [83] considers as supertask hypercomputers the AUTM⁺ and relativistic hypercomputation. Regarding the latter, I have already argued that the PMH proposal is not physically realizable, due to its not satisfying the usability constraint. Thus, I turn my attention to the AUTM⁺. In short, the difference between the AUTM and the AUTM⁺ has to do with the inclusion of a second moment limit stage in the latter such that one of the machine’s stages is not completely determined by the configuration of the previous stage (unlike a Turing machine). The reason, as Shagrir clarifies, is that “there are infinitely many stages between any given stage that precedes the second moment and the second moment limit stage” [83, pg. 54]. Hence, this not only violates Gandy’s Principle I, but is also not physically realizable, as it violates the fifth sub-constraint due to the Argument by Infinity.

Copeland and Shagrir [19], in a separate paper, do agree with Shagrir’s conclusions in [82] that the AUTM is not a hypercomputer, but assert that this depends on how one views computation and a Turing machine. In particular, they mention that an AUTM can be a hypercomputer (specifically, it can compute the halting function) if one is a “Turing-machine realist” and views computation as “external computation”. A Turing-machine realist, in their terminology, is someone who believes that Turing machines are not located within the ontology of pure mathematics (namely, as n -tuples whose members are sets and functions, such as the well-known Lewis-Papadimitriou definition of a Turing machine as a quadruple), but are notional machines. The level of notional machines lies between pure-mathematical ontology and the physically-realizable level, adhering to principles of temporality and causality. The authors then proceed to cite examples from Turing’s original paper [100] that appear to support that he was a “realist”. Nevertheless, this is all circumstantial evidence, for as even the authors admit – while Turing’s descriptions of his machines indicate temporality (such as the machine being “set in motion” [100, pg. 232], for instance), he provides no *explicit* discussion of the temporal aspects of Turing machines. But this is not a very important matter. What is important is what Copeland and Shagrir mean by “external computation” because it is only when it coincides with realism does the AUTM appear to be a hypercomputer. As the name suggests, external computation involves an appeal to the behavior of a device or system external to the machine (as opposed to internal

computation, which prohibits communication with the outside world during computation). Furthermore, unlike internal computation, which Copeland and Shagrir mention is consistent with Turing’s original definition of computation in [100], external computation “involves an extension to the classical theory of Turing machines, but one that is fully consistent with the realist’s view of the Turing machine” [19, pg. 236]. It is interesting to note the apparent inconsistency here. While Copeland and Shagrir were content to *circumstantially* justify their notion of Turing machine realism based on excerpts from Turing’s original work, they are not at odds with *explicitly breaking* Turing’s original definition of computation when adhering to the notion of external computation. The Turing machine is a model of internal computation, and this has been accepted since 1936, beginning with Turing himself! Furthermore, the authors fail to distinguish between the notions of absolute and relative computability. Had they, then they would have even more clearly seen (and in a precise manner) that the idea of external computation is incredibly alien to the concept of the Turing machine; interactive computation (which is an example of relative computability) commonly exploits external computation, whereas the Turing machine belongs to the notions of absolute computability. The distinction between absolute and relative computability and its relevance to computer science practice and hypercomputation will be discussed in detail in §4.

3.2.5. *Comparative notions of deterministic machines.* Moreover, Copeland and Shagrir [18] argue that PMH is of interest because it *prima facie* complies with Gandy’s Principles I-IV, yet computes the non-computable. Namely, PMH is a discrete state machine that consists of two communicating digital computers and PMH is consistent with the laws of general relativity, yet there are two important differences which Copeland and Shagrir are keen to point out. First, Gandy assumes that processes that consist of infinitely many steps do not terminate, whereas PMH allows terminating process that consist of infinitely many steps. Second, if one adheres to Gandy’s preliminary definition of deterministic that “the subsequent behaviour of the device is uniquely determined once a complete description of its initial state is given” [39, pg. 126], then PMH is deterministic in this sense. But, the more mathematically precise definition of what Gandy means by “deterministic” presented in Principle I explicitly requires that the configuration of each state is uniquely determined by the configuration of the previous step; PMH does not follow Principle I, and thus is not deterministic in Gandy’s sense. Copeland and Shagrir argue that PMH can be considered deterministic in a sense that is not “Gandy-deterministic” based on the initial state of the machine, and that “this sense of determinism is in good accord with the physical usage whereby a system or machine is said to be deterministic if it obeys laws that invoke no random or stochastic elements” [18, pg. 228]. However, if the PMH proposal is unlikely to be physically realizable (since it does not comply with the usability constraint), then Gandy’s definition of determinism cannot be justifiably ruled as too narrow on the grounds of not being “in good accord with the physical usage”.

3.2.6. *Analog hypercomputation.* Perhaps what is more interesting is the physical possibility of analog hypercomputation (which Copeland and Shagrir also address), since as mentioned at the beginning of §2, the discussion of physical properties (which can be used to possibly construct an analog hypercomputer) that cannot be simulated by Turing machine has spanned over at least four decades, even before the “hypercomputation movement” began. I have hitherto discussed in this section the impossibility of discrete deterministic hypercomputers existing in the real world based on Gandy’s principles, and have ignored the issue of analog hypercomputers for the time being, primarily due to the fact that Gandy did not discuss the analog case in his original paper [39]. Such proposals for instance that I have excluded from consideration include those by Kreisel [54][55], Pour-El and Richards [78][79], and da Costa and Doria [21]. Due to a lack of space, I will not go into the details of these proposals, especially since several of them (in particular the Pour-El and Richards results) have been well-known for quite some time. Instead, my purpose is to formulate a response to the physical realizability of analog hypercomputation in general. In fact, in an unpublished hand-written manuscript, Gandy [40] approaches analog machines (which were excluded in Thesis M) and discusses examples of physical systems (including the aforementioned proposals), both classical and quantum mechanical, for devices that compute the non-computable. His argument is that given a reasonable definition of “analog

machine”, such a machine cannot calculate the values of a number-theoretic non-computable function, upon accepting a continuously variable computable input. Gandy’s position is summarized as follows (cf. Kieu [48, pp. 557-8] and Copeland and Shagrir [18, pp. 221-2]):

If we let A be a recursively enumerable non-recursive set (for example the set which represents the halting problem), then a decision problem can be posed as the question of whether or not an integer $j \in A$. Because it is recursively enumerable, there exists a total computable function $a : \mathbb{N} \rightarrow \mathbb{N}$ which enumerates A without repetitions. Next, define the *waiting-time* function ν as

$$\nu(j) := \mu n[a(n) = j],$$

which is a partial recursive function whose domain is A , and which is not bounded by any total computable function. Moreover, $\nu(j)$ gives the least n (“time”) upon which j is confirmed to be a member of A . Now, Gandy argues that for any particular analog machine there is an upper bound J on the inputs it can accept, and he then proceeds to define the total function $\beta(J)$ as

$$\beta(J) := \max\{\nu(j) : j < J \text{ \& } j \in A\}$$

(with $\max \emptyset = 0$). $\beta(J)$ is not computable, and indeed it eventually majorizes every computable function. Gandy’s claim is that with given J , one cannot design an analog machine (whose behavior is governed by standard physical laws) which will give correct answers to all the questions $j \in A$? for $j < J$ unless one knows a bound B for $\beta(J)$. To illustrate the significance of the wording of the claim, suppose (what is quite plausible) that someone proves that $j \notin A$ for all $j < J = 10$; then he can design a machine which always outputs “NO” for $j < J$. But, because of this proof he does in fact know that $\beta(J) = 0$. Of course, if one knows a bound B as above, then one does not need an analog machine to settle the decision problem. One simply computes $a(n)$, for all $n < B$, to see if j is obtained as a value, and thus belongs to A , or not.

It is worth mentioning that Gandy’s argument, though rather encompassing, was intended as a “challenge rather than a dogmatic assertion” [18, pg. 221]. Kieu attempts to take on this challenge by arguing that Gandy’s conception of an analog machine is too narrow, stating

Gandy also recognised that his claim might not stand up to the kind of analogue machines whose behaviour depends on a single quantum (similar to our process involving Hamiltonians which have discrete spectra of integer values). Neither would the claim be valid if the physical theory involved has elements of non-computability already built-in. But here we have argued all along that Quantum Mechanics with its measurement postulate has, in a sense, some elements of non-computability through its *intrinsic randomness*. To wit, randomness is ‘noncomputable’ as it cannot be algorithmically computed....Furthermore, our proposal is not subject to Gandy’s treatment since ours is a kind of probabilistic computation, where the probability that we get a wrong answer can be made arbitrarily small but is not exactly zero. Also, his upper bound B turns out to be a product of our quantum hypercomputation: it is an output of the computation rather than a required input. Kieu [48, pg. 558]

It is interesting further to mention that Copeland and Shagrir mention Kieu’s counterarguments to Gandy’s analog machine claim, writing that “we will not pursue these matters here, but observe that the deep question of whether there are (ideal) analogue machines which satisfy usability constraints and which, when provided with a computable input, will give a non-computable output, is open” [18, pg. 222]. However, certainly Kieu’s proposed hypercomputer does not satisfy the usability constraint we are using, for it certainly does not satisfy the fifth sub-constraint, due to the Argument by Infinity as mentioned in §2.3 (and if one is to accept Kieu’s proposal based on occupation numbers then another sub-constraint, namely the first sub-constraint, is violated since the output may not be readable). For now, it is a very strong possibility that quantum mechanics (the weakness which Kieu alleges Gandy’s claim to be susceptible to) is unlikely to allow for physically realizable hypercomputation (even despite the fact that, Kieu’s proposal aside, quantum operations call for the execution of primitive operations that Turing machines cannot perform); instead, a more interesting prospect is that it may provide great future aid in solving problems of current practical

intractability, which I will discuss in §4.2.1.

3.2.7. Gandy’s *reductio ad absurdum*. In sum, as can be seen, unlike with Turing’s and Gandy’s definitions of computability, for which one can obtain a precise axiomatization, there is currently no single set of principles under which all hypercomputational models (provably) fall. Of course, even if there is no single set of principles, all it would take is one physically realizable notion of hypercomputation that goes beyond the universal Turing machine to justify this idea; nonetheless, given the usability constraint, the arguments in this section have supported Gandy’s principles as sufficiently general for real machines. Instead, one attains some degree of mathematical precision in terms of understanding the question “What is hypercomputation?” as a weakening or contravention of at least one of Gandy’s four principles (at least for non-analog hypercomputers, the analog case was dealt with separately). Regarding the issue that Copeland and Shagrir take with Gandy’s *reductio ad absurdum* argument, namely, as Israel puts it, that each of Gandy’s principles is necessary to avoid the “absurdity or vacuity...[that] every number-theoretic function is computable” [47, pg. 197], they assert that each number-theoretic function *is* computable, relative to some set of capacities and resources. As a result, they do not agree that there is any absurdity in allowing that every number-theoretic function is computable (relative to some index), especially since questions about which functions are computable relative to certain physical theories are seldom trivial. “When classicists say that some functions are absolutely uncomputable, what they mean is that some functions are not computable relative to the capacities and resources of a standard Turing machine. That particular index is of paramount interest when the topic is computation by effective procedures. In the wider study of what is computable, other indices are of importance” [18, pg. 229]. But, my arguments did not appeal to the capacities and resources of a standard Turing machine since the gedankenexperiment of the Turing machine was not originally intended for physical computation but regarded the limits of effective procedures. Instead, my assertions were based on the epistemic usability constraint which is reasonable given current, modern physical theories. Thus, Gandy’s axiomatic principles for mechanisms appear to be general enough to make hypercomputation physically unrealizable, and for now, given the arguments presented in this section, it appears that what Piccinini calls the Modest Physical CT⁸ holds, namely, that any function that is physically computable is (Turing-)computable, for which Piccinini already provides extensive support.

3.3. The “test” for hypercomputation: black box approach. Another way to obtain a general negative answer regarding the physical realizability of hypercomputation is not through a purely axiomatic approach to the problem, but by also treating a hypercomputer as a “black box”. And instead of being concerned with the question of *how* the device computes the non-computable (as I was heretofore), I shall only be concerned with *whether* the devices compute the non-computable. In other words, only the input and outputs of a “black box” hypercomputer will be considered, and we will try to understand whether one can adequately produce a test for hypercomputation, without any regard to the particular inner-workings of models.

For the purposes of this section, the “black box”, despite its name, may be characterized in a mathematically precise manner. Leitsch, Schachner, and Svozil [56] do precisely this, and I will only provide here a brief recapitulation of their results in order to avoid too much repetition. They first define the following⁹

Definition 1 (Preliminaries). Define B to be a subset of $X_1 \times \dots \times X_m$. Now, define the i -th projection of B with $i = 1, \dots, m$, as

$$B_i := \{x | x \in X_i, (\exists y \in X_1 \times \dots \times X_{i-1})(\exists z \in X_{i+1} \times \dots \times X_m)(y, x, z) \in B\}.$$

Furthermore define for any $x \in \mathbb{N}^m$ the following

$$|x| := \max\{x_i | i \in \{1, \dots, m\}\}.$$

⁸“Modest” with respect to the Bold Physical CT, which states that *any* physical process is (Turing-)computable.

⁹Note that I will use the same notation as in [56] so as to avoid any confusion.

Definition 2 (Black box). Let X, Y be sets. A subset of B is defined as a black box if $B_1 = X$, where X is the input set and Y is the output set.

Definition 3. Let B be a black box, define

$$f_B := \{(x, y) | (\exists z)(x, y, z) \in B\},$$

$$t_B := \{(x, z) | (\exists y)(x, y, z) \in B\}.$$

f_B is the input/output relation of B and t_B is the computing time of B . If f_B and t_B are functions, then B is, by definition, deterministic.

Definition 4 (Bound class). Define the bound class C to be a class of computable monotone functions $\mathbb{N} \rightarrow \mathbb{N}$ containing the polynomials over \mathbb{N} with nonnegative coefficients.

Definition 5 (Weak hypercomputer). Let M be a Turing machine computing a total function, $f_M : X \rightarrow Y$ to be the function computed by M , and let t_M be the computing time of M . Then

$$\{(x, f_M(x), t_M(x)) | x \in X\}$$

is called a black box. A weak hypercomputer is a black box B with the following property: There exists a bound class C such that

- $\forall M$ with $f_M = f_B$ and almost everywhere $\forall g \in C$, $t_M(x) > g(|x|)$.
- $\forall x \in B_1$, $\exists h \in C$ such that $t_B(x) \leq h(|x|)$.

Definition 6 (Strong hypercomputer). A strong hypercomputer is a black box B where f_B is non-computable.

The difference between a strong and weak hypercomputer is that a strong hypercomputer computes a non-computable function, whereas a weak hypercomputer outperforms all Turing machines. A possible scenario for a weak hypercomputer B is that if f_B is an EXPTIME-complete problem, then there exists no polynomial p and there exists no M such that $\forall x \in X$ it computes f_B with $t_M(x) \leq p(|x|)$, but $\forall x \in X$ and for a polynomial p , $t_B(x) \leq p(|x|)$. Furthermore, for stochastic hypercomputers, either t_B or both f_B and t_B are random variables, and the computational requirements have to be specified. For non-deterministic hypercomputers, one may distinguish between the following two cases, namely, f_B is not a function; or, f_B is a function, but t_B is not.

After establishing the above definitions, the authors consider a few cases and tests, and note that “these test methods will be essentially heuristic and present no way of systematically addressing the issue of falsifying or even verifying hypercomputation” [56, pg. 135]. I believe it to be a more pragmatic and reasonable approach to evaluate specific important cases and to prove that one cannot verify whether or not a hypercomputer can perform better than the known discrete computational models. The reason is that evaluating the operational validity of a black box hypercomputer for *all possible* cases is most likely rather unfeasible if only for the reason that there may be a large number of cases to consider, and the specific cases that are considered in [56] are ones in which the problems are considered easy in terms of their creation and hard in terms of their verification, and it is often times these “asymmetric” problems that hypercomputers are purported to solve (particularly in polynomial time) or decide. To avoid too much repetition, I shall describe only a couple of the scenarios that they consider, mainly to highlight the important consequences of their black box interpretation. For instance, it is not possible to obtain a proof to verify that a black box hypercomputer can solve an NP-complete problem (such as the satisfiability problem) in polynomial time, particularly because one would need to select the appropriate problem sequences, not to mention that one would need to know at least some information about the polynomial in question, such as its maximum degree. Moreover, regarding the computation of random sequences (such as Chaitin’s Ω [8]) and quantum randomness to compute the non-computable, the authors demonstrate that “beyond heuristic testing, any general statement about quantum randomness remains conjectural” [56, pg. 140]. It is also proven that even if there is a strong hypercomputer that solves, for example, the halting problem, the verification of its correctness is unfeasible, and one cannot expect to obtain a proof, in this case, of halting or non-halting from the hypercomputer itself. The best humans could one day do would be to construct a series of tests (provided that the hypercomputer is not stochastic and wrong answers are admitted) to indicate behavior

structurally differing from computable proof systems – though this would still not verify hypercomputable properties of black boxes. The authors ultimately conclude that “if presented with a hypercomputer or oracle, we could only assert heuristic information, nothing more. We have to accept the fact that more general assertions, or even proofs for computational capacities beyond very limited finite computational capacities remain impossible, and will possibly remain so forever” [56, pp. 139-140].

As of now, the physical theories we have deem hypercomputation physically implausible. What may be required to build a hypercomputer is a groundbreaking new physical theory. And given that one cannot ever preclude the possibility that a hypercomputer is built thousands of years from now due to the possible advent of paradigm-shifting physical theories, the issue of testing whether or not a hypercomputer compute the non-computable becomes a severely relevant issue. Hence, we are faced with currently not just one, but at least two, incredibly difficult (and perhaps impossible) problems, namely building and subsequently testing these devices, as I have shown through two mathematically precise means via an axiomatic approach and via the black box interpretation of Leitsch, Schachner, and Svozil. Of course, in practice, when it comes to the latter issue of testing a hypercomputer, if one already knows the inner workings of the particular hypercomputer being tested, then testing it should not be an issue, but this presupposes the existence of some new physical theory which allows for the construction of the hypercomputer in the first place (unless of course the machine being tested is not a hypercomputer, which defeats the purpose). As another extreme case, which Davis [23, pg. 11] mentioned in 1958, one cannot exclude the possibility of being presented with a (perhaps extremely complex) device that computes a non-computable function (possibly by extraterrestrial visitors). In this case, an understanding of the inner workings of the device may not be known beforehand or even permitted to be known, so thus, testing does become problematic. Now, it is important to mention that I am not saying that the theoretical pursuit of hypercomputation and the models presented so far are to be discarded, since some may have interesting theoretical applications (e.g. the *o*-machine which was seminal to the notion of uniform relative computability as well as the PMH proposal which is an intriguing contribution to debates on the foundations of physics); however, in the case of physical realizability, these models are not to be considered for that purpose. Therefore, to make hypercomputation relevant to the problems and applications of today, a shift in emphasis must occur from computing the non-computable to expressing/modeling certain processes in a superior manner than the closed-system paradigm of the Turing machine.

4. THE PRACTICAL LIMITATIONS OF TURING-EQUIVALENT MACHINES: A PROPOSAL

4.1. Computational theory vs. Computational practice. While hypercomputation has been primarily concerned with making the theoretical limits of computing potentially physically realizable, complexity issues in computational practice have not been given much focus (if at all) by this movement. As an immediate important example, which Davis [26] brings to light, the work of Cook and Levin on NP-completeness made severely relevant the issue of problems for which feasible algorithms are not known and for which Turing computability does not help, since the number of steps required to solve the problem grows exponentially with the length of the input, making their use in practice a problem. Certainly the lofty, to say the least, goal of computing the non-computable in practice is far from the already difficult problem of solving intractable problems in practice.

As Feferman [38] distinguishes: The aim of computational practice is to produce hardware and software that is, among other things, reliable, efficient, and user-friendly. The aim of the theory of computation, through its employment of logic and mathematics, is to provide limits to what is feasible and to provide a body of concepts around which to organize experience and a body of results predicting correctness, efficiency, and versatility in order to aid engineers in producing hardware and software that meet the aforementioned requirements. Certainly this comparison coincides with the usability constraint, which I used in §3 to evaluate the possibility of physical hypercomputation more generally. While the focus of this paper so far has either implicitly or explicitly been on the universal Turing machine, its use in computational practice is questionable, and Turing-equivalent models are often more applicable to a certain area than Turing’s *a*-machine, particularly because the Turing machine was not intended to be a practical computing device. Nonetheless, there still remain limitations in computational practice for which Turing computability does not

help, and this will underly my proposal for a more plausible hypercomputability. And in order to explicate and justify my assertions when examining the practical limitations of Turing-equivalent machines, I will appeal to the usability constraint.

Many machines that are thought to have a smaller instruction set, use less memory, or are more applicable to a certain area to solve problems than a Turing machine, can be shown not to have more computational capacity than the Turing machine [46]. It is important to note that Turing’s *a*-machine was not meant to model a “real” computer, but to determine the limits of effective procedures, and of course, a computer with fixed internal storage was developed only *after* Turing’s influential paper. I will provide here a few limited examples relevant to briefly illuminate this concept. For instance, Wang [104] introduces a variety of Turing-equivalent machines, but most relevant to this discussion is his basic machine, or B-machine, a Turing-equivalent computational model with only 4 sequential instructions (as opposed to the 7 sequential instruction of the Post-Turing machine¹⁰), and as Minsky mentions, is “the first formulation of a Turing-machine theory in terms of computer-like models” [63, pg. 200]. In fact, it is well-accepted that the Turing-machine itself does not provide a realistic model of actual computers and that register machines more closely resemble modern digital computers [93, §3.5]. However, when it comes to the question of the importance of the Church-Turing thesis *in practice*, not just for hardware but for software as well, Feferman [38, pp. 338-9] cites the following response from Odifreddi:

- (1) The notion of the universal Turing machine is the idealization and conceptual precursor of modern all-purpose computers.
- (2) The Enumeration Theorem shows the equivalence of programs and data, which is basic for stored program machines.
- (3) The proof of Kleene’s Normal Form Theorem for partial recursive functions provides the underlying theory for interpreters.
- (4) Various definitions of recursiveness provide the computational core and style of different programming languages such as PASCAL and LISP.

We are in agreement with Feferman that nevertheless, given the above points about the importance of absolute computability to practice, the notions of relative computability have a much greater significance in the realm of computational practice than those of absolute computability. It is here that Turing’s *o*-machine (which certainly is not a “forgotten idea” of Alan Turing, as Copeland and Proudfoot [14] claim) becomes of prime relevance. Feferman attributes this to the modular organization of hardware and software, citing the example of built-in functions in computers as “black boxes” (or “oracles”) in practice. He also mentions in a footnote that several people have suggested to him interactive computation as an example of Turing’s oracle in practice, but “while I agree that the comparison is apt, I don’t see how to state the relationship more precisely” [38, pg. 340, fn. 8]. Given that this was written in 1992, and since then technology has greatly expanded and is now significantly centered around the World Wide Web, online computation is the *most* pervasive example of Turing’s oracle in day-to-day computational use. And, in fact, in a 2009 article, Soare [94] does provide a precise characterization of the relationship between Turing’s oracle and online computing. It is important to note that the Turing machine itself does not communicate with its environment (i.e. whatever is external to it) during computation. But, on the other hand, online processes do involve interaction with the environment, not to mention that the relatively commonplace notion of upgrading hardware or software on a computer offers new functionalities, and as a result, its programs change due to external influence. This is problematic for the Turing machine paradigm in terms of encapsulating such processes; therefore, the notions of absolute computability do not capture certain, rather ubiquitous, aspects of modern computing. Though, none of this requires one to go significantly beyond the theoretical framework of what was developed by the “classicists”, since

¹⁰The term “Post-Turing machine” is accredited to Martin Davis and is a Turing-equivalent model of computation comprising a variant of the model Post [76] had proposed independently of Turing in 1936, and though these models have their similarities, it is worth mentioning that Post’s paper contains neither the notion of a universal machine nor the undecidability results of Turing.

The theory of relative computability developed by Turing and Post and the *o*-machines provide a precise mathematical framework for database or online computing just as Turing *a*-machines provide one for offline computing processes such as batch processing. Oracle computing processes are those most often used in theoretical research and also in real world computing where a laptop computer may communicate with a database like the World Wide Web. Often the local computer is called the “client” and the remote device the “server”. Soare [94, pp. 370-1]

When it comes to computational practice, however, there are still important developments yet to be made regarding interactive computation that are not accounted for by the theoretical models, that I will discuss later on in this section.

4.2. Towards a plausible hypercomputability. Just as one refers to the “theory and practice of computability”, I will similarly refer to the “theory and practice of hypercomputability”. With regard to the *theory* of hypercomputability, I do not propose any drastic changes since I do not think that models of hypercomputation that supposedly compute the non-computable should be discarded since they may be of theoretical interest (not to mention that the theoretical study of non-computability has been an active field of research since the 1940s starting with Kleene and Post). On the other hand, concerning the *practice* of hypercomputability, it is the purpose of this final section to justify a new direction that hypercomputation should take in terms of its practical aims. Instead of attempting to provide some physical credence to computing the non-computable, hypercomputation should adequately encapsulate certain modern computational paradigms for which Turing computability does not help, such as those that are computationally intractable in practice, as well as interactive and online computing. §3, as I have already mentioned, has provided the impetus (in a precise manner) to shift the focus of hypercomputation, and I will proceed to specify my proposal for how this shift could occur below.

4.2.1. Intractability in practice. Even if a hypercomputer is built, there is still the question of the time it takes to provide answers. If there is significant delay, such as if the time it takes to produce an output is exponential with respect to the size of the input, then reliability (sixth sub-constraint) becomes an issue if the inputs it is given (for some particular purpose) are large enough that the output is never returned in a finite amount of time that is useful for practical purposes (such as 1000 years, 10^{10} years, etc.). However, one does not need hypercomputation to realize that complexity is a relevant issue in the real world – these issues are a well-known occurrence in current computational practice. The question that arises is, how can one effectively solve certain important problems that are intractable? A major step forward would be to change the computational paradigm used in practice from the commonplace electronics. Quantum computing, especially, may help to overcome calculations that are computationally infeasible (not to mention, as alluded to in §3.2.6, that not all of the primitive operations of a quantum Turing machine, introduced by Deutsch [29], can be performed by a person simply using pencil and paper – but it should be noted that it cannot compute non-recursive functions). Already this is underway, the most famous example being Shor’s algorithm [84] in 1994, which demonstrated that the factoring of large integers could occur in polynomial time on a quantum computer, a task insurmountable for the standard electronic computer (and particularly relevant to cryptography, for example). Quantum computers that are able to do this have yet to be built; hence, the fifth sub-constraint of physical constructibility is not yet satisfied, though their realizability is certainly not impossible (an excellent discussion is provided in Nielsen and Chuang [68, Ch. 7]). The field of quantum computing has many opportunities for new innovations, and thus one goal of hypercomputation could be to make such quantum computations physically realizable. Despite this initial optimism, there is also evidence regarding the limits of quantum computing (cf. Nielsen and Chuang [68, Ch. 6.7]). This is based on a black box approach, namely, there is a boolean function f that takes an n -bit string as an input which can be evaluated at any input of our choice, and we are to determine if $\forall x, f(x) = 0$ or if $\exists x$ such that $f(x) = 1$. The end result is that in computing boolean functions using a black box, quantum algorithms

may provide only a polynomial speedup over classical algorithms *at best* and even this is not possible in general. No exponential speedup over classical algorithms is possible without additional information about the *structure* of the black box, which means that this does not rule out, for instance, a quantum polynomial time algorithm for an NP-complete problem such as 3-SAT. However, this cannot come just from understanding the quantum model, but from an algorithmic breakthrough on the specific problem itself. My mention of quantum computing is intentionally brief as its potential is rather well publicized, and I will discuss some other lesser known alternatives to electronics that show interesting potential in the realm of tractability.

Introduced by Păun [72], P systems (which were mentioned in §2.4), abstract distributed parallel computing models from the structure and function of a living cell. There are many variants of P systems, but the basic model of P systems is a membrane structure consisting of several cell membranes which are hierarchically embedded in a main membrane, termed the skin membrane. Membranes delimit regions, wherein objects are placed. Objects evolve according to evolution rules. These rules are applied to objects in the region with which they are associated and can modify the objects, send them outside the current membrane (or to an inner membrane), and can dissolve the membrane. The whole functioning of the system is governed by evolution rules which are applied in a maximally parallel manner, namely, at each step, all objects that can evolve should evolve. A *computation* starts from an initial configuration and proceeds using the evolution rules, and the computation is considered completed when it “halts” – no further rule can be applied to the object present in the last configuration. P systems are both capable of Turing universal computations and, in cases where enhanced parallelism is provided, can solve intractable problems in polynomial time. As an example of the latter, Păun [73] proposes a class of P systems where evolution rules are associated with both objects and membranes, and membranes can not only be dissolved, but can multiply themselves by division. In this way, the number of membranes can grow exponentially, and by making use of this enhanced parallelism, Păun proves that one can compute faster, namely, the well known NP-complete SAT problem can be solved in polynomial (linear) time (more on the computational efficiency of P systems can be found in a survey article of Păun [74, §8]). The computational universality of P systems and the possibility of solving NP-complete problems in linear time makes membrane computing a promising candidate as a model of computation. But it is important to mention that there it is still an open question of implementing a P system of this type in either biochemical or electronic media; therefore, there is still work to be done regarding whether or not this P system can fully satisfy the fifth sub-constraint of physical constructibility. An informative article briefly overviewing the fast occurring developments in P systems is provided by Păun [74], and in §§9 and 10 of this article, he discusses current research topics in the study of P systems as well as their applications.

Computational mathematics may be feasible with DNA (beyond just primitive mathematical operations). The key advantage to DNA is that it can make computers much smaller than before, while at the same time maintaining the capacity to store prodigious amounts of data. Since Adleman’s [1] pioneering paper, DNA computing has become a rapidly evolving field with its primary focus on developing DNA algorithms for NP-complete problems. Unlike quantum computing in recent years, though, the viability of computational mathematics on a DNA computer has not yet been fully demonstrated. In fact, only recently have the primitive operations in mathematics (i.e. addition, subtraction, multiplication, and division) been fully implemented and optimized. A recent paper of mine [66] discusses the theoretical possibility of implementing fast matrix multiplication algorithms (Strassen’s algorithm is the focus¹¹) with DNA due to the massive parallelism offered by DNA, and fast matrix multiplication serves as a bottleneck for many algorithms (such as inversion, computing determinants, and graph theory). Of course, there is no current advantage that DNA has over the electronics with regards to mathematical computing, not to mention the question of determining reliability [57], and at the moment DNA computing has merged to controlling and mediating information processing for nano structures and molecular movements. It is also relevant to note that Sieg [86] mentions that the Gandy machine is an abstract mathematical definition that embodies generally accepted properties

¹¹There have been several recent improvements to the exponent on the theoretical runtime of matrix multiplication (by Stothers [97] and Williams [109]), which is conjectured to optimally be $O(n^2)$ for the multiplication of two $n \times n$ matrices. But unlike Strassen’s algorithm [98], these algorithms are not used for matrices of practical sizes ($n < 10^{20}$, [66]).

of parallel computations (and indeed, with John Byrnes, he formulates an abstract model for parallel computations in [85]), and that they may be simplified via a graph-theoretic or a category-theoretic representation. However, “what is needed most...is their further mathematical investigation, e.g., for issues of complexity and speed-up, and their use in significant applications, e.g., for the analysis for DNA computations or of parallel distributed processes” [86, pg. 402]. Of course, whether or not Gandy machines are the most “favorable” model of parallel computations is a separate matter, but Sieg’s consideration does bring to light the question of complexity issues in DNA computation. Hence, another goal of hypercomputation could be to address the issues concerning reliability with DNA computations (which is relevant to the sixth sub-constraint of the usability constraint), as well as understanding with mathematical rigor an investigation of complexity issues that arise in DNA computations. And an overall long-term goal would be to pursue other computational paradigms that would allow for intractable problems in practice to potentially become tractable. As a starting point towards this long-term aim, Eberbach and Wegner [34, §5] consider four guiding principles with which to deal with intractability (and I include the paradigms considered so far here in this list), namely:

- (1) **Interaction:** through feedback with the external world, where the “world” can be represented by an environment or other computational agents, and “feedback” can be in the form of direct advice or by some performance measure obtained by interaction with the environment describing how good a solution is and/or the resources used to find it. Performance measures are used by evolutionary algorithms, anytime algorithms, and the $\$$ -calculus (“cost calculus”), for example.
- (2) **Evolution:** transforming a problem to a less complex or incomplete one. This approach is used by structural programming, approximation algorithms, anytime algorithms, P systems, and the $\$$ -calculus, to name a few.
- (3) **Guessing:** selecting randomly a path to find a solution. For instance, this approach is used by nondeterministic, probabilistic, ergodic, and evolutionary algorithms.
- (4) **Parallelism:** exploring all possible solutions simultaneously usually by trading time complexity for space complexity. This approach is employed by quantum computing, DNA computing, P systems, and the $\$$ -calculus.

Based on these principles, they present three other alternatives than the ones mentioned here for dealing with intractability: evolutionary computation, anytime algorithms, and the $\$$ -calculus. The reader is referred to the relevant sections (§§6.3-6.5) of [34] for a detailed discussion of these alternatives. Hence, not only are there six broad proposals for making hard computational problems tractable in practice, with each proposal offering many opportunities for further research, but the principles of interaction, evolution, guessing, and parallelism may also be used to serve as the basis of further inquiry into seeking out new methods of successfully dealing with intractability.

4.2.2. Paradigms of interactive computation. The above mention of parallel computations and interaction allows for the perfect segue into discussing the role hypercomputation could play in interactive computing. Before doing that, a few misconceptions must first be avoided. Eberbach, Goldin, and Wegner [35] define the following:

Super-Turing Computation. *All computation, including that which cannot be carried out by a Turing machine.*

And they further mention three principles that allow them to derive “super-Turing” models of computation:

- (1) Interaction with the world;
- (2) Infinity of resources;
- (3) Evolution of the system.

One specific example of what they call “super-Turing computation” is the dynamic interaction of clients and servers on the Internet because the “resulting concurrent interaction of multiple clients and servers, with a dynamic configuration of communication links and nodes, cannot be described as a Turing Machine computation, which must be sequential, static, and with all input predefined” [35, pg. 178]. Their use of

“super-Turing computation” is misleading, as it could be misconstrued as allowing for hypercomputation and that given their example of clients and servers, this every-day interaction (in the modern world) could provide a means of physically implementable hypercomputation. But there is a difference between *expressing/encapsulating* processes in a manner superior to that of the Turing machine, and computing more than a Turing machine. While certainly (as I have previously established) the “closed-system” paradigm of the *a*-machine does not allow for the encapsulation of interactive processes, it must be made clear that the three principles of “interaction with the world”, “infinity of resources”, and “evolution of the system” do not allow for physically implementable hypercomputation. In fact, in §6.1 of their paper, they present several ways of solving the halting problem via the principles of interaction, infinity, or evolution. However, they remark that “none of these solutions of the halting problem (*H*) are algorithmic. In the proofs, either some steps of the solutions do not have well defined (and implementable) meaning – this is when we refer to the help of oracles – or we use infinite resources (an infinite number of steps requiring perhaps an infinite initialization time)” [35, pg. 189]. Thus, it is clear that none of these “hypercomputations” are physically realizable, either due to the Argument by Infinity or the Argument by Physical Vagueness.

It is further interesting that they mention Turing’s work relating to the *o*-machine. In fact, a subsequent inquiry into this aspect of Turing’s research does reveal a framework for online computing long before the advent of the World Wide Web, therefore making Goldin and Wegner’s bold statement rather unnecessary:

Turing machines had been accepted as a principal paradigm of complete computation, and it was premature to openly challenge this view in the late 1970s and the early 1980s....A paradigm shift is necessary in our notion of computational problem solving, so it can provide a complete model for the services of today’s computing systems and software agents. Goldin and Wegner [41, pg. 102]

What Wegner, Goldin, and Eberbach fail to clarify is the distinction between absolute and relative computability. Clearly, the notions of absolute computability (such as the Turing machine and the Church-Turing thesis) do not properly encapsulate interactive processes. None of this requires a paradigm shift from the theoretical groundwork laid by the “classicists”, for the mathematically precise notions of relative computability become relevant here. Soare [94, pg. 382] combines the work of Turing [101, §4] and Post [77, §11] on relative computability and restates it in modern terms as follows

Post-Turing Thesis. *A set B is effectively reducible to another set A iff B is Turing reducible to A by a Turing oracle machine ($B \leq_T A$).*

Although technically speaking, the distinction between absolute and relative computability does not really require a separate thesis, Soare’s presentation does explicate this distinction rather nicely. He notes that while Turing’s brief introduction of oracles in his 1939 doctoral dissertation [101] did not state this as a formal thesis, it is largely implied by his presentation. Furthermore, Post [77] in 1944 makes this thesis explicit, claiming that it is the formal equivalent of the notion of *effectively reducible*, “a step as significant as the Church-Turing characterization of “calculable”” [94, pg. 382]. Soare further asserts that under the Post-Turing thesis, real world online or interactive processes can be described by Turing’s *o*-machine, and that “discussing *only* Turing *a*-machines in modern texts, or *only* the Church-Turing Thesis, and not the Post-Turing Thesis on oracle computers, is like discussing only batch processing machines of the 1950’s long after the emergence of online computing” [94, pg. 388]. Hence, though more needs to be mentioned in the computer science literature about relative computability (since an asymmetrical emphasis has been placed on the results of absolute computability), it already offers a rigorous mathematical framework for interaction, unlike Wegner’s interaction machines [106].

However, the important take-away from all this is that nevertheless, the notion of Turing computability does not suffice for adequately encapsulating interaction. As an example to explicitly illustrate this in relation to modern computing systems, I will consider the Actor model, a concurrent model of computation first introduced in 1973 by Hewitt, Bishop, and Steiger [43] (subsequently developed by Clinger [12] and Agha [2], among others), one of a number of message-passing models developed in the 1970s, along with Hoare’s

Communicating Sequential Processes (CSP) and Milner’s π -calculus. Conventional models of sequential computation (such as the Turing machine, the λ -calculus, etc.) employ the notions of *global time* and *global state*, wherein computation is carried out from one well-defined state to the next, and the transitions between global states are linearly ordered in the global time of the system. This global state approach was continued in automata theory for finite state machines and push down stack machines (as well as for their nondeterministic versions). When computation is not sequential, these notions may not be appropriate. Instead, certain concurrent systems (such as multiprocessor systems and geographically distributed systems [12]) are better analyzed by splitting the global state into pieces and viewing the net computation as a set of local computations that interact through message passing. The Actor model emphasizes these notions of local time and local state, and in contrast to the boundedness and locality constraints for Turing computers, the Actor model does not carry out computation in one place, from one well-defined state to the next; computation is distributed in space, and “actors” are treated as universal primitives of concurrent (digital) computation. Local times are represented by the arrival orderings of actors, which operate independently of each other except when they interact through message passing, and the communications between actors are represented by the activation ordering. From the actor point of view, sequential computations are distinguishable from concurrent computations by their event diagrams, and are in fact a special case of concurrent computations in that the activation ordering of a sequential computation is linear and its event diagram is a conventional execution sequence.

The Actor model is also of use in computational practice – an example of its industry uses are in the programming languages Erlang and Scala. Erlang was first developed by Joe Armstrong in 1986, designed to improve the development of telephony applications [3]. Designed for concurrency, distribution, and scalability, actors are a part of Erlang itself. Scala (which stands for “scalable language”) is a more recent programming language that runs on the Java platform and interoperates with all Java libraries, implementing actors using the multi-threaded capabilities of the Java Virtual Machine (JVM), while removing the complexity of thread communication for programmers [69] [64]. In fact, the popular social-networking site Twitter is using the Actor model for scalability by rewriting their backend infrastructure from Ruby on Rails to Scala [64]. Moreover, the Open Systems Laboratory, led by Gul Agha, aims to simplify the development of scalable parallel, distributed, and mobile computing systems, basing their research on the Actor model [70].

Nevertheless, given that interaction is not adequately encapsulated by the Turing machine and the multitude of conceptual frameworks (e.g. CSP, the π -calculus, Actor model, etc.) for dealing with interaction, there are still certain important aspects of computational practice that are not well expressed by these frameworks. The underlying conclusion is that these frameworks are intended to provide an *abstract* framework for dealing with and reasoning about concurrency (and therefore abstract away from certain real world details). With the Actor model, for instance, timing details follow two general requirements. First, each actor always has the capability to process messages sent to it. Second, every message eventually arrives at its target, known as *finite delay*¹². These requirements leave much ambiguity about the order of events in an actor’s arrival ordering, and the resulting nondeterminism is called *arrival nondeterminism*. The reason for the ambiguity in these requirements is simple, namely that:

For a programming language semantics to specify completely the order in which events are to occur during multiprocessor execution of a concurrent program is generally impossible, since it would entail fixing myriad details such as the number and relative speeds of concurrent processors, the exact times of and delays occasioned by page faults and other interrupts, the

¹²In fact, the type of computation that usually involves actors and distributed process models does not relate to single algorithmic computations, but can be viewed as functional compositions of such algorithms, where despite the latency of their result (namely, the second property of finite delay), these interactions are oracles, no different from regular function calls. Latency is a rather natural property of abstract descriptions of concurrent systems; finite but unbounded delay gives rise to unbounded nondeterminism. Although the Actor model exhibits unbounded nondeterminism (due to the underlying property of fairness), it must be made clear that this does *not* mean that the Actor model can physically compute any functions outside the class of recursive functions.

timing of signals between processors and the manner in which they are arbitrated, and so on, down to the levels of time resolution at which quantum indeterminacy becomes important. Clinger [12, pp. 56-7]

However, in cyber-physical systems, there is a real need for theories that help in the design of these systems when not only interaction, but time plays a role; for example, at the device driver level, time is especially crucial. Time is not the only challenge that needs to be addressed. Networked Cyber-Physical Systems (NCPS) present many challenges not properly addressed by current distributed computing paradigms, such as limitations of computational, energy, and networking resources, as well as accounting for uncertainty, partial knowledge, and bad or stale information [67][96]. In relation to the usability constraint, NCPS are certainly “usable”, where an emphasis is placed particularly on the sub-constraints of repeatability, physical constructibility, and reliability. Although work on cyber-physical systems deals with the engineering problem of integrating software with the real world, the problem lies deeper than that as “it is becoming increasingly clear that more fundamental work is needed to address real-world challenges in a uniform and systematic way” [67]. The work of Stehr, et al. [67][96] aims to explore a new notion of software that behaves closer to a physical or biological system. This is motivated by the fact that in challenging environments, the physical world imposes severe limitations on how distributed algorithms can operate, and instead of viewing resource limitations as a bug, they view it as a feature – and traditional models of distributed computing are too abstract to provide a proper foundation. The applications of NCPS are multifarious, providing complex, situation-aware services for traffic control, medical systems, automated laboratories, deep space exploration, emergency response, or social networking, to name a few. In fact, biological systems provide many examples of an effective NCPS, such as the human immune system, which as examples of desirable features has robustness, distributed knowledge, and generic and adaptive responses to events [96, pg. 112]. There is much more that can be said about NCPS, so I refer the reader to Stehr, et al. [96] instead, where a long-term vision as well as research directions pertaining to NCPS are contained. The development and applications of NCPS offer an exciting venue for hypercomputation in terms of visiting the theoretical shortcomings of distributed computing paradigms and their adaptations to the physical world. While the abstract notions of relative computability have been around since the 1940s and were further greatly developed in the 1970s and 1980s, offering a mathematical framework for interactive and online computing that the α -machine could not adequately provide, they are not sufficient to address the many challenges of real world computing, and this is not just an engineering problem. What is needed in distributed computing is a closer synergistic relationship between foundational concerns as well as a consideration of the real world details that are crucially relevant to computational practice.

5. CONCLUSIONS

The hypercomputation movement, in its current form, is certainly not of interest to contemporary computer science practice. Since certain models of hypercomputation may hold considerable import to the theoretical aspects of computer science (as well as to debates on the foundations of physics, such as the PMH proposal), I am not saying that these models should be discarded altogether either. Rather, my interest stems from practical concerns regarding hypercomputability. Given that the hypercomputation movement is rather young, I believe that a shift in the focus of hypercomputation is for the best. Instead of asserting that non-computable functions, tasks, etc. are physically computable, a more reasonable field of hypercomputation should be focused on modeling computational processes that are not adequately encapsulated by the Turing machine (such as contemporary problems dealing with certain issues of interaction and tractability) – the question of whether or not these processes allow one to compute the non-computable is not considered. But why should my proposal be considered? Through a rigorous analysis of Turing computability and by adhering to rather general and reasonable epistemic constraints on physical computation, I demonstrated that Gandy’s axiomatic principles for machine computation are sufficiently general given current physical theories. And further, there is no physical theory even remotely on the horizon that could not only permit the construction of a hypercomputer, but adequately produce a test for hypercomputation, when its inner

workings are not known or cannot be known. Considering that hypercomputation is focused on surpassing the limits on computation imposed by absolute computability (namely, computing more than the universal Turing machine), it is doubtful that even its goal would be of much interest to current computer science practice. Questions regarding issues of complexity as well as the notions of relative computability are examples of questions of prime interest in computational practice. In particular, my proposal for a shift in hypercomputation contained issues for which Turing computability does not help, namely having to do with exploring alternative computational paradigms for effectively dealing with intractable problems in practice as well as building robust cyber-physical systems, and the scope of these issues are general enough for there to be ample opportunity for further research, yet specific enough for their importance to computational practice to be well-known and accepted. It is these aforementioned problems that a plausible field of hypercomputability should focus on in order to be practically viable.

ACKNOWLEDGEMENTS

I am grateful to Professor Solomon Feferman for his invaluable guidance throughout the process of this research. I also express my gratitude to Steven Ericsson-Zenith for discussions on interactive computing and the historical premises of Turing's work, and to Carolyn L. Talcott for discussions on the Actor model and cyber-physical systems. I further thank Martin Davis and Wilfried Sieg for their helpful comments on a draft of this article.

REFERENCES

1. Adleman, Leonard. "Molecular computation of solutions to combinatorial problems". *Science* **266** (1994): 1021-1024.
2. Agha, Gul A. "Actors: a model of concurrent computation". Doctoral Dissertation. MIT Press. 1986.
3. Armstrong, Joe. "History of Erlang". In *HOPL III: Proceedings of the third ACM SIGPLAN conference on History of programming languages*. 2007.
4. Barwise, J., H. J. Keisler, and K. Kunen (eds.). *The Kleene symposium*. Amsterdam: North-Holland, 1980.
5. Blake, Ralph M. "The paradox of temporal process". *Journal of Philosophy* **23** (1926): 645-654.
6. Calude, Cristian S. and Gheorghe Păun. "Bio-steps beyond Turing". *Biosystems* **77** (2004): 175-194.
7. Calude, Cristian S. and Boris Pavlov. "Coins, quantum measurements, and Turing's barrier". *Quantum Information Processing* **1** (2002): 107-127.
8. Chaitin, Gregory J. "A theory of program size formally identical to information theory". *Journal of the ACM* **22** (1975): 329-340.
9. Church, Alonzo. "An unsolvable problem of elementary number theory". *American Journal of Mathematics* **58** (1936): 345-363.
10. Church, Alonzo. "Review of Turing 1936-1937". *Journal of Symbolic Logic* **2** (1937): 42-43.
11. Church, Alonzo. "On the concept of a random sequence". *Bulletin of the American Mathematical Society* **46** (1940): 130-135.
12. Clinger, William. "Foundations of actor semantics". Doctoral Dissertation. MIT Press. 1981.
13. Copeland, B. Jack. "Turing's o-machines, Searle, Penrose and the brain". *Analysis* **58** (1998): 128-138.
14. Copeland, B. Jack and Diane Proudfoot. "Alan Turing's forgotten ideas in computer science". *Scientific American* **253** (1999): 98-103.
15. Copeland, B. Jack and Richard Sylvan. "Beyond the universal Turing machine". *Australasian Journal of Philosophy* **77** (1999): 46-66.
16. Copeland, B. Jack. "Accelerating Turing machines". *Minds and Machines* **12** (2002): 281-300.
17. Copeland, B. Jack. "Hypercomputation". *Minds and Machines* **12** (2002): 461-502.
18. Copeland, B. Jack and Oron Shagrir. "Physical computation: how general are Gandy's principles for mechanisms?". *Minds and Machines* **17** (2007): 217-231.
19. Copeland, B. Jack and Oron Shagrir. "Do accelerating Turing machines compute the uncomputable?". *Minds and Machines* **21** (2011): 221-239.
20. Costa, José F., Pedro Rodrigues, and Siegelmann, Hava T. "Verifying properties of neural networks". *Proceedings of the 6th International Work-Conference on Artificial and Natural Neural Networks: Connectionist Models of Neurons, Learning Processes and Artificial Intelligence-Part I* (2001): 158-165.
21. da Costa, Newton C. A. and Francisco A. Doria. "Classical physics and Penrose's thesis". *Foundations of Physics Letters* **4** (1991): 363-374.
22. Davies, E. Brian. "Building infinite machines". *British Journal for the Philosophy of Science* **52** (2001): 671-682.
23. Davis, Martin. *Computability and unsolvability*. New York: McGraw-Hill, 1958.
24. Davis, Martin, Hilary Putnam, and Julia Robinson. "The decision problem for exponential Diophantine equations". *Annals of Mathematics* **47** (1961): 425-436.
25. Davis, Martin. "The myth of hypercomputation". In [99]: 195-211.
26. Davis, Martin. "The Church-Turing thesis: consensus and opposition". Lecture notes in computer science, volume 3988 (2006): 125-132.
27. Davis, Martin. "Why there is no such discipline as hypercomputation". *Applied Mathematics and Computation* **178** (2006): 4-7.
28. de Leeuw, K., E. F. Moore, C. E. Shannon, and N. Shapiro. "Computability by probabilistic machines". In [60]: 183-212.
29. Deutsch, David. "Quantum theory, the Church-Turing principle, and the universal quantum computer". *Proceedings of the Royal Society, Series A* **400** (1985): 97-117.
30. Dijkstra, Edsger W. *A discipline of programming*. Prentice-Hall, 1977.
31. Earman, John and John D. Norton. "Forever is a day: supertasks in Pitowsky and Malament-Hogarth spacetimes". *Philosophy of Science* **5** (1993): 22-42.
32. Earman, John. *Bangs, crunches, whimpers, and shrieks: singularities and acausalities in relativistic spacetimes*. Oxford University Press, Oxford, 1995.
33. Earman, John and John D. Norton. "Infinite pains: the trouble with supertasks". In [65]: 231-261.
34. Eberbach, Eugene and Peter Wegner. "Beyond Turing machines". *Bulletin of the European Association for Theoretical Computer Science* **81** (2003): 279-304.
35. Eberbach, Eugene, Dina Q. Goldin, and Peter Wegner. "Turing's ideas and models of computation". In [99]: 159-194.
36. Echeverría, Javier, Andoni Ibarra, and Thomas Mormann (editors). *The space of mathematics: philosophical, epistemological, and historical explorations*. Walter de Gruyter, Berlin, 1992.
37. Etesi, Gábor and István Németi. "Non-Turing computations via Malament-Hogarth space-times". *International Journal of Theoretical Physics* **41** (2002): 341-370.
38. Feferman, Solomon. "Turing's 'oracle': from absolute to relative computability—and back". In [36]: 314-348.

39. Gandy, Robin O. "Church's thesis and principles for mechanisms". In [4]: 123-148. 1980.
40. Gandy, Robin O. "On the impossibility of using analogue machines to calculate non-computable functions". Unpublished. 1993.
41. Goldin, Dina Q. and Peter Wegner. "Computation beyond Turing machines". *Communications of the ACM* **46** (2003): 100-102.
42. Hagar, Amit and Alex Korolev. "Quantum hypercomputation - hype or computation?". *Philosophy of Science* **74** (2007): 347-363.
43. Hewitt, Carl, Peter Bishop, Richard Steiger. "A universal modular Actor formalism for artificial intelligence". *Proceedings of the 3rd international joint conference on Artificial intelligence* (1973): 235-245.
44. Hodges, Andrew. "[FOM] on Kieu from Andrew Hodges". May 2004. <http://www.cs.nyu.edu/pipermail/fom/2004-May/008129.html>.
45. Hogarth, Mark. "Does general relativity allow an observer to view an eternity in a finite time?". *Foundations of Physics Letters* **5** (1992): 73-81.
46. Hopcroft, John E. and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation*. Addison-Wesley Publishing Company: Reading, Massachusetts: 1979.
47. Israel, David. "Reflections on Gödel's and Gandy's reflections on Turing's thesis". *Minds and Machines* **12** (2002): 181-201.
48. Kieu, Tien D. "Quantum hypercomputation". *Minds and Machines* **12** (2002): 541-561.
49. Kieu, Tien D. "Quantum algorithm for Hilbert's tenth problem". *International Journal of Theoretical Physics* **42** (2003): 1461-1478.
50. Kieu, Tien D. "Quantum adiabatic algorithm for Hilbert's tenth problem: I. The algorithm". October 2003. <http://arxiv.org/abs/quant-ph/0310052>.
51. Kieu, Tien D. "Reply to Andrew Hodges". March 2006. <http://arxiv.org/abs/quant-ph/0602214>.
52. Kleene, Stephen C. *Introduction to metamathematics*. Amsterdam: North-Holland, 1952.
53. Kleene, Stephen C. *Mathematical logic*. New York: Wiley, 1967.
54. Kreisel, Georg. "A notion of mechanistic theory". *Synthese* **29** (1974): 11-26.
55. Kreisel, Georg. "Review of Pour-El and Richards". *Journal of Symbolic Logic* **47** (1982): 900-902.
56. Leitsch, Alexander, Günter Schachner, and Karl Svozil. "How to acknowledge hypercomputation?". *Complex Systems* **18** (2008): 131-143.
57. Losseva, Elena. "Building reliable systems from unreliable components and the DNA computing paradigm". *Bulletin of the European Association for Theoretical Computer Science* **81** (2003): 305-317.
58. Löwe, B., A. Sorbi, and S. B. Cooper. (eds.) *New Computational Paradigms: Changing Conceptions of What is Computable*. Springer, Heidelberg, 2007.
59. Matijasevič, Juri V. "Diofantovost' perechislimykh mnozhestv". *Doklady Akademii Nauk SSSR* **191** (1970): 279-282 (Russian). English translation in: Matijasevič, Juri V. "Enumerable sets are Diophantine". *Soviet Mathematics. Doklady* **11** (1970): 354-358.
60. McCarthy, John and Claude E. Shannon (editors). *Automata studies*. Princeton University Press, Princeton, New Jersey, 1956.
61. McCulloch, Warren S. and Walter H. Pitts. "A logical calculus of the ideas immanent in nervous activity". *Bulletin of Mathematical Biophysics* **7** (1943): 115-133.
62. Mead, Carver and Lynn Conway. *An introduction to VLSI systems*. Addison-Wesley, Reading, Massachusetts, 1980.
63. Minsky, Marvin. *Computation: finite and infinite machines*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1967.
64. Mok, Waiming. "How Twitter is scaling". <https://waimingmok.wordpress.com/2009/06/27/how-twitter-is-scaling/>. 27 June 2009.
65. Morton, Adam and Stephen P. Stich (editors). *Benacerraf and his critics*. Cambridge, Massachusetts: Wiley-Blackwell, 1996.
66. Nayebi, Aran. "Fast matrix multiplication techniques based on the Adleman-Lipton model". *International Journal of Computer Engineering Research* **3** (2012): 10-19. Also available at: <http://arxiv.org/abs/0912.0750>. December 2011.
67. "Networked Cyber-Physical Systems at SRI". <http://ncps.csl.sri.com/>.
68. Nielsen, Michael A. and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2010.
69. Odersky, Martin, Lex Spoon, and Bill Venners. "Scala: a scalable language". <http://www.artima.com/scalazine/articles/scalable-language.html>. 6 May 2008.
70. "Open Systems Laboratory: Research". <http://osl.cs.uiuc.edu/research.php>.
71. Ord, Toby. "Hypercomputation: computing more than the Turing machine". Honours thesis, The University of Melbourne. September 2002. <http://arxiv.org/abs/math/0209332>.
72. Păun, Gheorghe. "Computing with membranes". *Journal of Computer and System Sciences* **61** (2000): 108-143.
73. Păun, Gheorghe. "P systems with active membranes: attacking NP-complete problems". *Journal of Automata, Languages and Combinatorics* **6** (2001): 75-90.
74. Păun, Gheorghe. "Membrane computing". In: *Fundamentals of Computation Theory*. Lecture Notes in Computer Science, Volume 2751 (2003): 177-220.
75. Piccinini, Gualtiero. "The physical Church-Turing thesis: modest or bold?" *British Journal for the Philosophy of Science* **62** (2011): 733-769.

76. Post, Emil L. "Finite combinatory processes. Formulation I". *Journal of Symbolic Logic* **1** (1936): 103-105.
77. Post, Emil L. "Recursively enumerable sets of positive integers and their decision problems". *Bulletin of the American Mathematical Society* **50** (1944): 284-316.
78. Pour-El, Marian B. and Jonathan I. Richards. "A computable ordinary differential equation which possesses no computable solution". *Annals of Mathematical Logic* **17** (1979): 61-90.
79. Pour-El, Marian B. and Jonathan I. Richards. *Computability in analysis and physics*. Berlin: Springer, 1989.
80. Russell, Bertrand A. W. "The limits of empiricism". *Proceedings of the Aristotelian Society* **36** (1936): 131-150.
81. Shagrir, Oron and Itamar Pitowsky. "Physical hypercomputation and the Church-Turing thesis". *Minds and Machines* **13** (2003): 87-101.
82. Shagrir, Oron. "Super-tasks, accelerating Turing machines and uncomputability". *Theoretical Computer Science* **317** (2004): 105-114.
83. Shagrir, Oron. "Supertasks do not increase computational power". *Natural Computing* **11** (2012): 51-58.
84. Shor, Peter W. "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer". *SIAM Journal on Computing* **26** (1997): 1484-1509.
85. Sieg, Wilfried and John Byrnes. "An abstract model for parallel computations: Gandy's thesis". *The Monist* **82** (1999): 150-164.
86. Sieg, Wilfried. "Calculations by man & machines: conceptual analysis". In [87]: 396-415.
87. Sieg, Wilfried, Richard Sommer, and Carolyn Talcott (editors). *Reflections on the Foundations of Mathematics*. Association for Symbolic Logic, 2002.
88. Sieg, Wilfried. "Calculations by man & machine: mathematical presentation. *Proceedings of the Cracow international congress of logic, methodology and philosophy of science. Synthese Series*, Kluwer Academic Publishers, 2002: 245-260.
89. Sieg, Wilfried. "Church without dogma - axioms for computability". In [58]: 18-44.
90. Siegelmann, Hava T. "Computation beyond the Turing limit". *Science* **268** (1995): 545-548.
91. Siegelmann, Hava T. *Neural networks and analog computation: beyond the Turing limit*. Boston, Massachusetts: Birkhäuser, 1999.
92. Smith, Warren D. "Three counterexamples refuting Kieu's plan for "quantum adiabatic hypercomputation"; and some uncomputable quantum mechanical tasks". *Applied Mathematics and Computation* **178** (2006): 184-193.
93. Soare, Robert I. "Computability and recursion". *Bulletin of Symbolic Logic* **2** (1996): 284-321.
94. Soare, Robert I. "Turing oracle machines, online computing, and three displacements in computability theory". *Annals of Pure and Applied Logic* **160** (2009): 368-399.
95. Stannett, Mike. "The case for hypercomputation". *Applied Mathematics and Computation* **178** (2006): 8-24.
96. Stehr, Mark-Oliver, Carolyn Talcott, John Rushby, Pat Lincoln, Minyoung Kim, Steven Cheung, and Andy Poggio. "Fractionated software for Networked Cyber-Physical Systems: research directions and long-term vision". *Lecture Notes in Computer Science, Formal Modeling: Actors, Open Systems, Biological Systems* **7000** (2011): 110-143.
97. Stothers, Andrew. "On the complexity of matrix multiplication". Ph.D. thesis, University of Edinburgh. <http://www.maths.ed.ac.uk/pg/thesis/stothers.pdf>. 2010.
98. Strassen, Volker. "Gaussian elimination is not optimal". *Numerische Mathematik* **13** (1969): 354-356.
99. Teuscher, Christof (editor). *Alan Turing: life and legacy of a great thinker*. Springer, Berlin, 2004.
100. Turing, Alan M. "On computable numbers, with an application to the *Entscheidungsproblem*". *Proceedings of the London Mathematical Society* **42** (1936-1937): 230-265. "On computable numbers, with an application to the *Entscheidungsproblem*. A correction". *Proceedings of the London Mathematical Society* **43** (1938): 544-546.
101. Turing, Alan M. "Systems of logic based on ordinals". *Proceedings of the London Mathematical Society* **45** (1939): 161-228.
102. Turing, Alan M. "Intelligent machinery". In [103]: 107-127. 1948.
103. Turing, Alan M. *Collected works of A. M. Turing: mechanical intelligence*. Darrel C. Ince (editor). Elsevier Science Publishers, 1992.
104. Wang, Hao. "A variant to Turing's theory of computing machines". *Journal of the ACM* **4** (1957): 63-92.
105. Wang, Hao. *From mathematics to philosophy*. London: Routledge and Kegan Paul, 1974.
106. Wegner, Peter. "Why interaction is more powerful than algorithms". *Communications of the ACM* **40** (1997): 81-91.
107. Welch, Philip D. "The extent of computation in Malament-Hogarth spacetimes". *British Journal for the Philosophy of Science* **59** (2008): 659-674.
108. Weyl, Hermann. *Philosophie der Mathematik und Natureissenschaft*. München; und Berlin: R. Oldenbourg, 1928 (German). English version in: Weyl, Hermann. *Philosophy of Mathematics and Natural Science*. Princeton University Press, Princeton, 1949.
109. Williams, V. Vassilevska "Breaking the Coppersmith-Winograd barrier". <http://www.cs.berkeley.edu/~virgi/matrixmult.pdf>. November 2011.

E-mail address: anayebi@stanford.edu

URL: <http://www.stanford.edu/~anayebi>